

DTIC FILE COPY

62

RADC-TR-90-340
Interim Report
December 1990



AD-A230 635

NORTHEAST PARALLEL ARCHITECTURES CENTER (NPAC) AT SYRACUSE UNIVERSITY

Northeast Parallel Architectures Center

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. 6263

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

DTIC
ELECTE
JAN 04 1991
S E D

Rome Air Development Center
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700

91 1 13 022

This report has been reviewed by the RADC Public Affairs Division (PA) and is releasable to the National Technical Information Services (NTIS) At NTIS it will be releasable to the general public, including foreign nations.

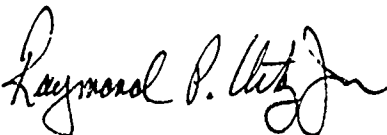
RADC-TR-90-340 has been reviewed and is approved for publication.

APPROVED:



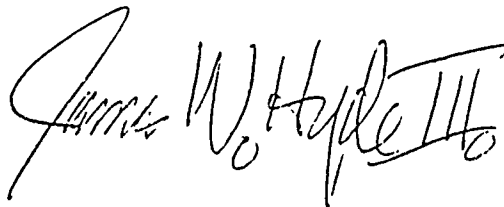
RAYMOND A. LIUZZI
Project Engineer

APPROVED:



RAYMOND P. URTZ, JR.
Technical Director
Directorate of Command & Control

FOR THE COMMANDER:



JAMES W. HYDE III
Directorate of Plans & Programs

If your address has changed or if you wish to be removed from the RADC mailing list, or if the addressee is no longer employed by your organization, please notify RADC (COES) Griffiss AFB NY 13441-5700. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

NORTHEAST PARALLEL ARCHITECTURES CENTER
(NPAC) AT SYRACUSE UNIVERSITY

Contractor: Northeast Parallel Architectures
Center
Contract Number: F30602-88-C-0031
Effective Date of Contract: 8 February 1988
Contract Expiration Date: 30 March 1991
Short Title of Work: Northeast Parallel Architectures
Center
Period of Work Covered: Jan 88 - Jan 89

Principal Investigator: William Schraeder
Phone: (315) 443-1723

RADC Project Engineer: Raymond A. Liuzzi
Phone: (315) 330-2925

Approved for public release; distribution unlimited.

This research was supported by the Defense Advanced
Research Projects Agency of the Department of Defense
and was monitored by Raymond A. Liuzzi (COES) Griffiss
AFB NY 13441-5700 under Contract F30602-88-C-0031.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE December 1990		3. REPORT TYPE AND DATES COVERED Interim Jan 88 - Jan 89
4. TITLE AND SUBTITLE NORTHEAST PARALLEL ARCHITECTURES CENTER (NPAC) AT SYRACUSE UNIVERSITY			5. FUNDING NUMBERS C - F30602-88-C-0031 PE - 62301E PR - F263 TA - 00 WU - 01	
6. AUTHOR(S)				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Northeast Parallel Architectures Center Syracuse University Syracuse NY 13214			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency 1400 Wilson Blvd Arlington VA 22209			10. SPONSORING/MONITORING AGENCY REPORT NUMBER RADC-TR-90-340	
			Rome Air Development Center (COE) Griffiss AFB NY 13441-5700	
11. SUPPLEMENTARY NOTES RADC Project Engineer: Raymond A. Liuzzi/COES/(315) 330-2925				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report summarizes what has been accomplished at the Northeast Parallel Architectures Center during its first full year of operation. This report is divided into two sections, Technical and Administrative, and represents an effort to measure the effectiveness of NPAC's mission: <p style="margin-left: 40px;">"The Northeast Parallel Architectures Center promotes parallel architectures and research support to university, corporate, and government researchers nationwide."</p>				
14. SUBJECT TERMS Computers, Parallel Computers, Architectures, Computer Hardware			15. NUMBER OF PAGES 122	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Technical Report Section I

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



1. Executive Summary

This one year report summarizes what has been accomplished at the Northeast Parallel Architectures Center during its first year of full operation. Under RADC contract, it encompasses all aspects of NPAC's Scope of Work as well as covers other important developments at the center. This report is divided into two sections, Technical and Administrative, and represents an effort to measure our own effectiveness in fulfillment of NPAC's mission:

The Northeast Parallel Architectures Center
promotes and explores advanced computing technology
by providing parallel architectures and research support
to university, corporate, and government
researchers nationwide.

2. Program Overview

The first year was a period to define the organization, acquire hardware and consolidate support and commitment to NPAC's mission. NPAC relied heavily on research faculty from other institutions and from Syracuse University (Colleges of Computer and Information Science, Engineering, Engineering (CASE)) to achieve this goal. This year was a period of rapid growth and development. The machines acquired include: two Connection Machines; two Multimaxes, a 320 and 310; an Alliant FX/80 and a Stellar Graphics Supercomputer. These architectures are discussed in further detail in part two of this report. We also continued to build internal and external relationships, strengthening our local and national reputation. Our facility matured and stabilized.

3. Contracts

Under this contract, NPAC provides high quality training and consulting for advanced parallel computers in several specific application areas. NPAC is expected to contribute significantly to advancing the national computational base in parallel computing. After discussions with DARPA and RADC, we recognized a clear need within DoD to build techniques and specific tools to assess parallel computer performance and we are working under a separate contract to develop such tools.

4. Development

A program plan was developed with input from RADC and DARPA, the University administration, and the Deans of CIS and Engineering. The plan required a fast-track build up in three phases over the past year: 1) acquire hardware and related infrastructure, 2) recruit staff, and 3) develop services required by RADC and DARPA. To select the architectures, RADC and DARPA program scientists and representatives from university faculty, administration, and staff were involved in the early phases of vendor negotiation. We recruited systems programmers to install and operate hardware and to begin the initial service, the Friendly User Program (FUP). During this interim service, NPAC invited University faculty, graduate students, and staff to use its computers, regardless of purpose, as long as they required little assistance. There were 760 users when the program was phased out in June 1988. The FUP allowed us to focus our attention on learning to operate the sophisticated and experimental computers, without overloading our limited personnel with user questions.

As a reliable national research center, NPAC established a strong administrative structure for long-term growth and to manage large sponsored research contracts and subcontracts. This is fully discussed in part two of this report.

NPAC is building visibility to facilitate broad and fast dissemination of advances and new knowledge in parallel computing. With planning and coordination, we anticipate that the flow of information will be two-way: out from NPAC and back from researchers. (The Technology Transfer Program is further discussed in part two of this report.)

5. NPAC Resources

5.1 Hardware

Hardware was acquired primarily under an Office of Naval Research grant to Syracuse University. Using commercially available hardware, NPAC offers both fine-grained and coarse-grained architectures that support multiple users at local and remote locations. The chart on the following page shows the current hardware configuration. The machines are explained in greater detail in part two of this report.

5.2 Software Tools

The development of software tools to guide or aid the user of parallel and/or distributed computing systems is a high priority within the community, and remains a relatively immature effort. NPAC's progress in this direction has, to date, been limited by competing demands. However, there are several significant tools installed on the systems. (In the brief descriptions that follow, please note that this discussion does not include those features and utilities supplied by the vendors as part of the operating system or user interface.)

X Window System, Version 11 (X11) is heralded by many to be the tool of the Nineties in computing. It is a portable, network transparent window system. Of particular interest is that the X Window server distributes user input to, and accepts output requests from various client programs through a variety of interprocess communication channels. Although the most common

situation is one in which all programs are running on the same computer, the X Window System can allow programs to be run transparently from other computers as well.

The Connection Machine Image Processing tool (CMIP) prototype, developed by Henry Krempel of the NPAC staff, is a mouse-oriented interface to Thinking Machines *Lisp system. It is intended to simplify and streamline the manipulation of 2-D image data. It allows users to load image data into the Connection Machine, filter the data and then display it. CMIP is built upon a variety of commercial products. Even as a prototype, it has demonstrated great potential value.

NPAC has constructed system monitoring tools which conveniently display computing activity on the various computers. While not intended to be a user tool at this time, some of this structure may be incorporated into a future tool to assist users in making scheduling and processor assignment decisions.

There are several tools that have been identified as potentially useful in the NPAC environment, some of which are not yet publicly available, which we are negotiating to obtain.

NPAC also supports *netlib*—a program to electronically distribute software for ALLUS, the Alliant User's Society. Netlib was developed by Eric Gross of Bell Labs and Jack Dongarra at Argonne National Laboratories. A number of user-contributed software packages are available through the system, such as various FORTRAN libraries and routines which were optimized for the Alliant, and system utility packages which were ported to run on the Alliant's architecture.

6. Networking

During the first year of NPAC operations under this contract, NPAC was the driving force behind these projects:

- The design and installation of a fiber optic based 10 Mbps Local Area Network (called FASTNET) which connects 12 buildings, several hundred researchers and several computing facilities at Syracuse University.
- The design and installation of several large SUN file servers, SUN workstations, and Ethernet LANs to provide a multi-tiered network linking to FASTNET as well as local support for workstations.
- Funded a study on how to best integrate a large number of high-performance workstations into the campus environment.
- Provided initial funding for an on-campus Workstation Integration Program which provides large discounts on workstations, installation, file servers, LANs, software, and instruction for researchers and educators.
- Organized and trained a strong cadre of UNIX users, then promoted and partially funded the acquisition of a campus run Gould NP-1 super-mini computer running UNIX.

Most of this equipment was purchased through other funding sources; NPAC established these programs to increase the impact of the funds provided under this contract.

In wide area networking, NPAC worked closely with NYSERNet to build stable and redundant T1 (1.5 Mbps) connections to the Internet for Syracuse University. NYSERNet is a regional WAN which accesses national networks as well as supercomputer and parallel processing facilities provided by NSF,

DARPA, NASA, and DOE. NYSErNet is part of Internet, comprised of nearly one thousand sites and one million computers.

7. Research Programs

In the broad field of advanced and parallel computing, applications range from the "hard" to the "soft" sciences. In recognition of this range, NPAC has established multiple research services to fully serve this complex mix of scientific needs and to give each individual scientist conducting research at NPAC the best support possible.

The Research Program service research areas which are likely to share a longer-term synergism with NPAC facilities and services. Through the selection of a small number of these program areas, NPAC can cultivate a variety of inter-project services to serve researchers' needs. These program areas will also foster new research opportunities that have a higher degree of promise for the scientific community in general, and NPAC in particular.

The list of program areas will be flexible, although not so dynamic as to change overnight, or with the current research buzzwords. The areas selected will reflect a careful assessment of those fields which are likely to have the largest impact on the use of parallel computers.

The current list of program areas is:

- computational neural science
- artificial intelligence
- signal processing
- software engineering
- algorithm development and evaluation

Each area will be supported by a full suite of services produced by NPAC for researcher use:

- Software libraries, with free distribution when possible

- Quarterly technical series with abstracts
- Active outreach
- Regularly-scheduled workshops and seminars
- Domain-specific user groups
- Regularly updated literary bibliographic search
- An annual meeting of national stature in at least one Research area

Our first annual meeting, in the field of Computational Neural Science, will be held in early autumn, 1989.

7.1 Research Profiles

Research in a wide variety of areas is currently being pursued using the resources NPAC provides. Projects granted allocations of time on the parallel computers range from nuclear magnetic resonance, molecular dynamics and parallel language development to real-time imaging, vision, oil field exploration and quantum chromodynamics. Scientists are employing the parallel architectures to implement new algorithms and achieve speedups in processing time. They are also investigating the application of parallel computers to established codes and methodology to solve ongoing scientific inquiries.

Categorical characterizations and descriptions of research approaches are prepared for each research project granted use of the NPAC resources in Research Profiles, snapshots of the work in progress. This collection of information is routinely and widely made available to demonstrate the applicability of parallel computing across the range of research topics.

A Research Profile is created from the information supplied in the research proposal of each project. These reside in a database, and, with the full proposal submitted by the research team, enable the Research Support staff and administration of NPAC to set up opportunities for exchange among current and potential users. A complete set of Research Profiles follows on the next 74 pages:

Title: Three-dimensional reconstruction, manipulation and display of the brain to study the brain's control of blood pressure.

Research area: Image analysis

Principal Investigator: Scott B. Berger, Cornell University Medical College

Architecture and Language: Connection Machine, *Lisp

For more than 30 years, we have studied the brain's control of blood pressure through a wide variety of experimental approaches, including anatomical, physiological and biochemical methods. One of their image analysis tools is a facility for 3D reconstruction of the brain, from serial histological sections, or from any other imaging modality which yields 2D consecutive sections (CT, MRI, PET, etc.). Our projects require a true volumetric representation, which is a formidable problem for computation and storage, and for which we designed a system based on a hierarchical data structure. This design is suitable for large scale parallel machines, and we have written a set of *Lisp algorithms which perform 3D reconstruction, manipulation and display of ultra-high resolution data sets.

Research area: Computational fluid dynamics

Title: Computational algorithms for solving unsteady nonlinear problems in fluid dynamics

Principal Investigator: Daniel T. Valentine, Clarkson University

Architecture and Language: Alliant FX/80, Fortran

The aims of the proposed project are: (1) To develop high-order, vectorized computational algorithms that take advantage of parallel architectures to solve unsteady nonlinear problems in fluid dynamics and, in particular, flow problems that are modelled by the Navier-Stokes equations. (2) To apply the computational methods to predict the development of flow separation over axisymmetric, afterbody geometries. (3) To extend the method to handle afterbody flows with shear-layer control (i.e. with suction and/or blowing at the flow boundaries). (4) To pay particular attention to post processing time-dependent flow predictions (i.e. to apply computer graphics capabilities to produce movies of the time-dependent solutions in order to study the

development of flow separation, the transients associated with an instantaneous imposition of a shear-layer control device, among other unsteady flow problems.

Title: Modeling a small neural network
Research area: Vision
Principal Investigator: Robert B. Barlow, Jr., Syracuse University
Architecture: Connection Machine

I am interested in applying the capabilities of parallel architectures toward understanding how the brain works. In particular, my interest is in the function of the relatively simple neural network of the retina of the horseshoe crab, which has about 1000 receptor units that integrate visual information with a combination of excitory and inhibitory mechanisms. This is the only neural network for which there exists a quantificative formulation based on physiological measurements. Parallel architectures will permit us to model the time-dependent properties of the eye in real or near real time and thus compare the model results with direct physiological recordings. Computer simulation will allow us to test the consequences of the brain's modulation of the retinal function. The retina changes state from day to night, and we are anxious to learn the relative importance of each component of the state change.

Research area: Tools for parallel program development
Title: Design, evaluation and tuning of application programs on parallel architectures
Principal Investigator: Hany H. Ammar, Clarkson University
Architectures: Encore Multimax and the Connection Machine
The decreasing cost of Supermini multiprocessor systems has made parallel processing technology available to a wide range of institutions, companies, and government agencies. Tuning application programs to the underlying parallel architecture is essential to exploit the full advantage of parallel processing. In this project, we propose the development of a trace driven

simulator that facilitates the tuning process of application programs. It also helps in evaluating the performance of the architecture and the detection of system bottlenecks which aids in upgrading the system by adding more memory or more CPU's. It is planned to apply this technique to image processing algorithms on the Connection Machine and to numerical algorithms on the Encore Multimax.

Research area: Computer graphics, image synthesis

Title: Massively parallel algorithms for computer graphics

Principal Investigator: Donald P. Greenberg, Cornell University; Peter Kochevar, Cornell University

Architecture and Language: Connection Machine, C/Paris

This work investigates the possibilities that massively parallel architectures may provide for the generation of realistic image synthesis. This research entails developing highly distributive algorithms for modeling shape and form, and for simulating light's interaction with matter. These algorithms will be implemented on the connection machine by assigning processors to regions of space and then arraying them into a three-dimensional grid with the stipulation that adjacent processors in the lattice correspond to proximate regions of space. Light paths will be constrained to follow lattice links and the sum over all paths from light sources to each lattice site will be computed inductively by all processors in parallel. The memory contents of certain designated measurement sites will collectively constitute the desired image at the end of a simulation.

Title: Synchronization of pacemaker activity within the sinoatrial node

Research area: Human Physiology

Principal Investigator: Donald C. Michaels, SUNY Health Science Center

Architecture and Language: Connection Machine, C*

In our current research, mathematical models are being used to investigate synchronization of pacemaker activity within the sinoatrial node (the natural

pacemaker of the heart) and propagation of electrical activity within and between various cardiac tissues. Cells are coupled through ohmic resistances to form one-, two-, and three-dimensional arrays. We have begun applying the tools of non-linear dynamics and chaos theory to the analysis of the patterns of activity obtained in our simulations. These tools allow us to investigate the complex and irregular dynamics observed under various conditions. We would like to test the feasibility of parallelizing existing Fortran code for this project. In particular, the two-dimensional arrays of cells should map nicely to the architecture of the Connection Machine, since each cell can be assigned to a single processor, and the interactions between cells can be modeled with the transfer of information between processors.

Research area: Numerical analysis

Title: Research on Alopex optimization algorithms

Principal Investigators: Erich M. Harth, Theodore Kalogeropoulos, Volker Weiss, Syracuse University

Architectures and Languages: Alliant FX/80, Fortran; Connection Machine, C*

Optimization of a physical situation generally means finding the "best" value of a function F , called the cost function, within a specified domain of adjustable variables. The Alopex Method of optimization solves this problem for cost functions which may be non-linear in a large number of variables and provides solutions which avoid local extrema by the stochastic introduction of noise. Optimization by Alopex algorithms has been successfully applied to a great variety of problems, including visual receptive fields, models of sensory perception, curve fitting, the traveling salesman problem, and crystal formation. Since the calculation of the algorithm for some of these problems involves large numbers of identical operations, the use of parallel computers could provide substantial speedups as well providing the capability of solving problems involving large numbers of variables, for example, many-particle systems.

Research area: VLSI design

Title: Simulate and formally verify VLSI circuits designed in MOS technology

Principal Investigators: Randal E. Bryant, Saul Kravitz, Derek Beatty, Carnegie Mellon University

Architecture and Languages: Connection Machine, C/Paris, C* and *Lisp

The COSMOS project seeks to achieve new levels of performance in simulating and formally verifying VLSI circuits designed in MOS technology. One of our aims is to exploit parallelism to achieve higher simulation performance. We plan to use both data parallelism, by simulating circuit operations over a number of input sequences independently, and circuit parallelism, by evaluating circuit elements in parallel while simulating a single input sequence. In addition, our project has developed symbolic simulation methods to formally verify that a circuit will implement its specification. Our Boolean manipulation algorithms create and manipulate complex graph data structures much like connectionist networks. Further research is required to determine the most effective ways to exploit the many sources of potential parallelism in these algorithms. Ultimately, we would like to integrate the Boolean manipulation code with our existing simulation code to implement a symbolic simulator.

Research area: Parallel language design and implementation

Title: Multiprocessor implementation of a parallel programming language

Principal Investigator: Per Brinch-Hansen, Syracuse University

Architecture: Encore Multimax

The proposed project is the first phase of a new research program at Syracuse University on programming methods for parallel architecture. The aim of the first project is to implement a parallel programming language called Joyce on the Encore Multimax computer and develop analytical models which accurately predict the performance of the implementation.

Joyce was initially designed and implemented by Per Brinch Hansen for IBM Personal Computers. A Joyce program defines a parallel computation in

which hundreds or thousands of processes are executed simultaneously. The processes communicate by exchanging messages. Joyce supports recursive processes and enables a compiler to perform complete type checking of messages. These features make Joyce a more powerful language than CSP and Occam, which are also message passing concurrent languages.

Research area: Computer vision

Title: To develop and implement parallel algorithms for computer vision

Principal Investigators: Hussein A. Ibrahim, Lisa G. Brown, Columbia University

Architecture and Language: Connection Machine, *Lisp

The goal of this project is to develop parallel algorithms in both low-level and medium-level vision. In the first year, several parallel stereo, depth interpolation, and texture algorithms for highly parallel computer architectures have been developed. They have also developed an environment to program pyramid and multi-resolution algorithms on the Connection Machine. Several of the parallel algorithms for stereo and texture have been implemented in this environment, including a new autocorrelation-based texture algorithm. In the pyramidal environment, an image is represented as an image pyramid consisting of a set of decreasing resolution versions of the image. Several classic pyramid operations have been used to test the environment such as stereo matching, edge refinement, pyramid search and similar logarithmically-improved pyramid operations. They propose to continue developing other parallel computer vision algorithms that use the multi-resolution approach to vision and that fuse information from cooperative processes to increase the certainty of computing surface parameters in the image.

Research area: Parallel programming languages

Title: A Connection Machine implementation of SUPER

Principal Investigators: J. Alan Robinson, Ernest Sibert, Syracuse University

Architecture: Connection Machine

Robinson and Greene proposed a high-level programming system known as SUPER (for Syracuse University Parallel Expression Reduction) which combines both functional and relational constructs. The language is defined by means of an abstract graph-reduction machine based on techniques pioneered by David Turner, but considerably extended to encompass the relational features of SUPER. Such a system naturally admits parallel reductions, without difficult synchronization requirements. The result is, at least in principle, a very expressive, generalized logic-programming system which can be given a parallel implementation, yet does not require programmers to specify parallel computations explicitly.

We have written a first implementation of SUPER on the Connection Machine, using a refined definition of the system. This effort has exposed a number of interesting questions concerning the abstract definition as well those raised by the limit of finite resources and those specific to the CM architecture.

Research area: Electromagnetics

Title: Develop a computer code capable of calculating multiple scattering effects from complex radar targets by taking advantage of parallel computing architectures

Principal Investigator: Paul R. Jedrezejewski, Syracuse Research Corporation

Architectures and Language: Connection Machine; Alliant FX/80, Fortran

The work proposed is to develop a practical radar cross section prediction scheme including multiple scattering effects using physical optics by taking advantage of parallel computer architecture. Physical optics is a high frequency scheme for predicting the radar cross section of a body. Depending on the direction of the incident field, the body is divided into an illuminated and a shadowed portion. The currents on the illuminated portion are

approximated by twice the component of the incident magnetic field intensity tangential to the surface. The currents on the shadowed portion are assumed to be zero. The currents are integrated over the surface of the scatterer resulting in the radiated far field. When the fields of interest lie in a specular direction (direction whose angle relative to the surface normal at the point of reflection is equal to the angle of incidence relative to the same normal), physical optics works very well. Multiply reflected fields may be observed at points of observation which are not specular to the incident radiation. The multiple scattering effects are what is of interest in this work.

Title: Solution of Partial Differential Equations via the
Finite Element Method - Connection Machine
Research area: Scientific Computation
Architecture: Connection Machine
Principle Investigator: Dr. Steven P Castillo, New Mexico State University

Currently, research is being conducted into the finite element solution of Laplace's equation, Poisson's equation, the wave equation, and the reduced wave equation. The solution of these differential equations and other large practical 2-D and 3-D problems is limited in size by time and memory constraints of current conventional supercomputers such as the CRAY XMP. Our project is to implement existing finite element code used for the solution of partial differential equations on the Connection Machine, addressing the issues: 1. How easy is it to convert or write code for the Connection Machine? 2. How fast is the Connection Machine for solving large problems in relation to conventional supercomputers such as the CRAY XMP? 3. How fast is the Connection Machine in comparison to other parallel architectures for solving large problems?

Title: Parallel Algorithms in Computational Geometry
Research area: Computational Geometry
Architecture: Connection Machine
Principle Investigator: Seoung-Jun Oh, Syracuse University

Computational Geometry is concerned with the design and analysis of computational algorithms for solving geometry problems. There are basically five categories in computational geometry according to the nature of geometric objects involved: convexity, intersection, geometric searching, proximity, and optimization. The framework for parallelism in computational geometry is efficient for both synthesis as well as analysis, since 1) all geometry problems can be described as collections of primitive geometry functions for which we can easily analyze the computational complexities, and 2) they are connected very systematically based on parallel problem solving techniques. Several algorithms in each primary area of computational geometry will be developed.

Title: Intelligent Spectral and Image Analysis
Research area: Nuclear Magnetic Resonance
Architecture: Fortran on Alliant, Connection Machine
Principle Investigator: George C Levy, Syracuse University

Analysis of nuclear magnetic resonance (NMR) of infrared spectra often requires reduction of spectral data near the limits of detection and is almost always accompanied by non-idealities such as undetermined baselines, incomplete feature separation, and inaccurate representation of feature line-shapes. Research is currently under way on three fronts: optimization of algorithms used for data processing in NMR and infrared spectroscopy, development of new algorithms for magnetic resonance imaging (MRI), and the combination of expert system and statistical analysis of spectra. Three programs for NMR and a program for infrared data processing have been

developed: NMR1 for analysis of 1D NMR spectra; NMR2 for 2D and 3D spectra and MRI; MEM for maximum entropy method processing and linear prediction; and SPECTIR for infrared data processing. These programs have several computationally intensive algorithms which can be speeded up by vectorization or by and by parallel processing of subregions of the spectral data.

Title: Study on Cursive Korean Character Recognition
using AI and Expert Systems
Research area: Artificial Intelligence
Architecture: C on the Multimax, *Lisp on the CM
Principle Investigator: Kyeung-Hak Seo, Syracuse University

Visual interpretation of cursive script is an interesting area in the fields of pattern recognition, image processing and A.I. Although there have been some studies of printed Korean character recognition, there has been very little work on cursive Korean character recognition. The proposed overall scheme for the character recognizer is that a set of knowledge-based experts are organized around a global data base, called a blackboard, which is the means of communication among the experts. The expert systems are run in parallel and consist of a scheduling expert, constraint propagation expert, stroke analyzing expert, loop finding expert, merging/splitting expert, character form decision expert, and a common sense expert to make choices among possible decisions.

Title: Analyze/Design microwave integrated circuits & develop novel approach for solution of operator equation in electromagnetics
Research area: microwave integrated circuits
Architecture: Alliant FX/80 (Fortran)
Principle Investigator: Guo Qing Wang, Syracuse University

Since computer circuits now have effective bandwidths of 10 GHz or higher, it is important to introduce propagation delay effects and dispersive effects in the analytical CAD models. In addition, the method of moments can be used in order to solve some problems in the electromagnetics. It is proposed that novel numerical approaches for the solution of equations with more than 1000 unknowns be studied. The Alliant FX/80 is perfectly suited to this kind of study.

Title: Develop system to monitor activity in I/O system of Connection Machine
Research area: Computer architecture
Architecture: Connection Machine
Principle Investigator: Junaid Zubairi,

Tools for monitoring I/O activity in a machine requires developing software to be loaded and run by the host computer simultaneously with the application software. This project will analyze the I/O performance of the Connection Machine. In particular, the performance of the message routing algorithms and router is to be evaluated. The system is to be loaded with a large number of messages and then, through Paris instructions, one can record the times that they arrive at their destinations. This has already enabled a Petri net model of the router to be built. Further evaluation of the routing algorithms will be done.

Title: Construct a logic simulation package and implement a simulated annealing algorithm
Research area: Computer Architecture
Architecture: Connection Machine
Principle Investigator: Songqing Cai, Syracuse University

I propose to build a logic simulation package that works at the gate level and to implement a simulated annealing algorithm for PLA logic minimization. The logic simulation package, which is for an independent study, simulates combinational and sequential synchronous circuits at the gate level and operates in an event-trigger manner to achieve maximal simulation speed.

Title: Load flow analysis on Parallel Computer Architectures
Architecture: Alliant, Multimax, Connection Machine
Principle Investigator: Dr. Kamal T Jabbour, Syracuse University

We propose to implement a load flow analysis program on a parallel architecture, to demonstrate the suitability of parallel processing to load flow analysis, and to evaluate performance improvement compared to traditional Von Neumann architectures. Since load flow and stability analyses are some of the most time-consuming computer tasks in power systems operations, we propose to study the structure of these algorithms, including both the Newton-Raphson and Gauss-Seidel techniques, for inherent parallelism. We will then identify the parallel computer best suited for this problem, write a parallel load flow analysis program, and evaluate its performance by comparison to a sequential program using the same technique.

Title: Examination of arithmetic algorithms in *Lisp for
the Connection Machine
Research area: Parallel numerical algorithms
Architecture: *Lisp on the Connection Machine
Principle Investigator: Patricia E Heywood, Syracuse University

This study involves an examination of and implementation of arithmetic algorithms for the Connection Machine, specifically, the 'classical algorithms', 'single precision calculations', and 'double precision calculations' found in Chapter 4 of Knuth's book, Seminumerical Algorithms. My work includes the development of parallel versions of Knuth's arithmetic algorithms for very large (or very small) integer and floating point numbers. Along with programming experiments on the CM in *Lisp to validate the algorithms, the problem of moving nonstandard numeric data from the Vax front end to the CM for this application will be explored.

Title: Parallel Graph Algorithms for the CM/PRAM
Simulation on the CM
Research area: Parallel Algorithms
Architecture: *Lisp on the CM1 and CM2
Principle Investigator: Todd Heywood, Syracuse University

This research consists of the investigation and implementation of published parallel graph algorithms which match a SIMD architecture well. The goal is to provide a collection of efficient parallel graph algorithms which will be useful for some common computationally intensive subproblems of many applications. Effort shall also be made to improve upon the known bounds of problems. It is also planned to implement an efficient simulation of the Parallel Random Access Machine (PRAM) on the Connection Machine. The

PRAM model can be used to express many parallel algorithms in an abstract way, without having to write details of communication, etc.

Title: Software Development for Molecular Dynamic
Simulation of Polymers
Research area: Molecular Dynamics
Architecture: Fortran on the Alliant and Multimax
Principle Investigator: Barry Hardy, Syracuse University

Molecular Dynamic simulations of polymer chains may afford the best possible means of understanding the mechanisms underlying polymer motions. The approach is based essentially on following the movement of the atoms of a model structure for sufficiently long time periods. The proposed research will follow this approach by: 1. The development and application of computer simulation methods to study polymer chain dynamics in the amorphous state. 2. The development of algorithms and analysis methods for the calculation of polymer chain statistics that can be compared with experiment. 3. The optimization and parallelization of these methods on a fast computer such as the NPAC Multimax or Alliant.

Title: Development of Parallel Program for the
Implementation of the Battle Management
Algorithm
Research area: Linear programming
Architecture: C on Encore Multimax
Principle Investigator: Shantanu Ganguly, Syracuse University

The Battle Management Algorithm is a computer system implementation of weapon to target allocation based on the relative weapon - target costs and availabilities. The problem is that of network linear programming, where the

weapon platforms are represented by supply nodes, the target platforms by demand nodes, and a cost is associated with each weapon - target arc. The Simplex Method of optimization has been chosen for the implementation of the algorithm. The problem matches closely with the 'Transportation Problem', and with certain modifications, the steps for its solution can be applied for the solution of the weapon allocation problem. Our aim is to implement both a serial and parallel version of the solution to this problem on the Encore Multimax. We plan to experiment with a number of test cases to minimize the run time and memory requirements of the program.

Title: Performance Analysis of LU Factorization on
Multiprocessor System
Research area: Analysis of Parallel Algorithms
Architecture: Multimax and Connection Machine
Principle Investigator: Mallick Shameem Ahsan, Syracuse University

The solution of a system of linear equations is an important computation found in many engineering and scientific applications. One commonly used approach to solve such a system is the LU factorization algorithm. Here, we propose a parallel decomposition of this algorithm and analyze its performance for hypercube multiprocessor systems and connection machines. Specifically, we consider the effects of problem decomposition, communication overhead and load balancing on the speedup.

Title: Development of Neural Net simulator
Research area: Neural Nets
Architecture: Connection Machine
Principle Investigator: Ganesh Ramamoorthy, Syracuse University

Connectionism refers to a class of massively parallel architectures that consists of several simple processing units, each connected to several other similar units. The knowledge in such systems is stored in the pattern of connections between the units and the corresponding weights. Since knowledge is stored by a pattern of activation over several processing units, it provides a distributed representation. Some connectionist models that are used as neural net classifiers are Hopfield net, Hamming net, Carpenter/Grossberg classifier, Kohonen's Self Organizing Feature Maps, and the Perceptron approach. Since a good deal of time is spent writing programs to emulate the networks under study, it would be immensely useful to have a framework which one could use for simulating all neural nets of this sort. It is proposed that such a neural net simulator be developed which would facilitate the simulations through network description.

Title:	Development of parallel algorithms for testing and simulation of logic circuits
Research area:	VLSI Design
Architecture:	C* and *Lisp on the Connection Machine
Principle Investigator:	Vinod Narayanan, Syracuse University

Parallel algorithms for testing and simulation of logic circuits will be developed. The algorithms will be intended primarily for SIMD architectures, and will be implemented on the Connection Machine. The specific problems addressed will be: 1) weighted random pattern testing based on test counts; 2) logic and fault simulation for gate level networks; 3) Switch level simulation; 4) Deterministic test generation. Solutions developed for these problems on non-parallel architectures are proving to be inadequate for the current complexity of VLSI designs. These problems have a considerable amount of data-level parallelism. The Connection Machine can result in significantly faster solutions to these problems. Such speed-ups will help ease the testing and design verification problems for large VLSI circuits.

Title: Modeling of Connection Machine and Parallel Architecture. Parallel algorithm analysis and development
Research area: Parallel Architectures and Algorithms
Architecture: C/Paris on the CM
Principle Investigator: Shan Jiang, Syracuse University

Parallel architectures give us opportunities to handle more complex, larger problems. At the same time, it makes the analysis of algorithms more complicated than that of using serial machines, which has been studied extensively. This research will work towards establishing a methodology to analyze parallel architectures and algorithms. It will define a model of the Connection Machine and develop a method to model other parallel architectures based on this experience. Some simulation tools will also be developed to analyze and compare some algorithms on the models.

Title: Conversion of serial loglisp to concurrent loglisp on Multimax
Research area: Computer Science
Architecture: Encore Multimax and Connection Machine
Principle Investigator: H R Nayak, Syracuse University

The basic objective of the research is the design of a parallel LOGLISP using concurrent searches. The existing sequential Lisp code will be converted for use with a concurrent Lisp. LOGLISP lends itself naturally to the application of parallel searches of the logic deduction tree, but ways must be found to minimize data sharing between multiple processes if significant parallel speed-ups are to be achieved.

Title: 1) Parallel algorithm for Programmable Logic Array
2) VLSI Circuit simulation
Research area: VLSI design
Architecture: Alliant and Connection Machine
Principle Investigator: Jim Kee Chan, Syracuse University

1) This project is to design and implement a parallel algorithm on the connection machine to implement Programmable Logic Array (PLA) logic optimization. We intend to first implement the heuristic method of implementing the general PLA algorithm, but we may also investigate graph-based algorithms and Boolean comparison. Performance will be compared with existing sequential implementations as well as other parallel ones. 2) A parallel sorter has been designed, and we need to validate the design by circuit simulation. Since the whole circuit contains two thousand transistors, the only practical circuit simulation tool is FX8spice on the Alliant FX/80. After the design has been validated, we will be fabricating a chip to implement the parallel sorter.

Title: Evaluation of algorithms for parallel computation of aerodynamic flows
Research area: Aerospace Engineering
Architecture: Fortran on Alliant FX/80
Principle Investigator: Mr. Edward A Bogucz, Syracuse University

External aerodynamic flows at large Reynolds numbers often are modeled in terms of two separate regions: an inviscid, irrotational flow away from the body, and a thin viscous boundary layer adjacent to the body surface. In this subdivision of the flow field, a different set of equations governs the flow in each region. The solution in each flow region must be matched through appropriate boundary conditions to obtain a composite representation of the

entire flow field. Practical calculation procedures based on this approach determine a solution of each flow field separately, and iterate until a consistent solution in both regions is obtained. The first part of this research is to obtain a more accurate description of inviscid flow fields with discrete vortices by developing parallel algorithms that use Lagrangian vortex methods rather than the usual Eulerian grid. The second part of this research will exploit concurrent solutions to the free-stream and boundary layer regions, and develop new coupling techniques that will improve the overall solution for flows in which the boundary layer has a strong effect on the free stream.

Title:	Exploring the use of parallelism in the development of constructive computational theories of social phenomenon
Research area:	Political Science
Architecture:	*Lisp on the CM, Lisp on the Multimax
Principle Investigator:	Stuart J Thorson, Syracuse University

We propose to continue work exploring the use of parallelism in the development of constructive computational theories of social phenomenon, along two fronts. The first involves the use of the Connection Machine to investigate properties of various social choice mechanisms. An example of a social choice problem is the election of a president. The basic idea behind these problems is that some sort of aggregation function is applied to a collection of individual preference orders to identify some 'best' social order. These sort of problems are typically investigated formally by showing that certain aggregation functions do or do not exist which meet certain conditions. The proofs are in general existence proofs (and not constructive proofs) perhaps in part because of the computational complexity involved. Thus our plan is to use the CM to 'reduce' the computational complexity by letting each processor contain the preferences of one agent. Our second line

of research has been to use a theorem proving program (ITP) to try to prove theorems about problems such as the one discussed above.

Title: Information Retrieval and Discourse Linguistics,
based on graph processing and connectionism.
Research area: Information Retrieval
Architecture: C* on the CM-2
Principle Investigator: Robert N Oddy, Syracuse University

The purpose of this project is to explore the application of discourse linguistics techniques to information retrieval from textual databases. We argue that information retrieval takes place in the context of problem-solving activities in the life of the user, and that keyword-based systems, even with semantic and statistical capabilities, have no way of capturing this context. The discourse-level structures of two important text types - user's problem statements and document abstracts - are seen to be closely related to the working situations of the users, and thus form suitable vehicles for an interactive retrieval system to negotiate with the user. We plan to automatically analyze empirical abstracts into discourse-level components, corresponding to aspects of empirical research methods. This information is incorporated into a connectionist model, where clue words and components are represented by nodes connected by weighted links (some inhibitory). Part of the network is fixed and represents general knowledge about the structure of abstracts, and part is specific to the features observed in the abstract under analysis. Following a period of interaction, activation levels of units are read off as indicators of the presence of components.

Title: Parallel & Vector Computing for Nonlinear Network Optimization
Research area: Mathematical programming
Architecture: C/Paris on CM-1 and CM-2.
Principle Investigator: Stavros A Zenios, The Wharton School

Nonlinear programming problems with network constraints arise in a wide range of engineering, management, statistical, economic and other applications. Such problems are typically characterized by their large size that may extend to thousands of variables and constraints. Our goal is to increase the size of solvable models and their domain of application via the design and implementation of suitable parallel algorithms and their testing with real data on parallel computers. Our proposed research is to design and implement distributed relaxation algorithms for a massively parallel computer. This involves a mapping of the network topology to the interconnection pattern of the computer and resolving issues like algorithm partitioning and synchronization. Emphasis will be places on communications, as well as computational, efficiency.

Title: Parallel Algorithms for Random Pattern Testability Computation
Research area: VLSI Design
Architecture: Connection Machine and Multimax
Principle Investigator: Pankaj Mayor, Syracuse University

SIMD and MIMD algorithms for computing random pattern testability in computer logic networks will be developed. The complexity of the algorithms will be determined analytically. The algorithms will be implemented on the Connection Machine and the Multimax. Performance of this algorithms will be compared with those that run on serial machines. Considerable speedup is

expected. The results from this work are expected to make a significant contribution to the field of computer logic testing.

Title: Implementation of alpha-L based on Dissertation
Research A Mathematics of Arrays
Research area: Parallel language design and implementation
Architecture: C and Fortran on Multimax and Alliant, C* on
Connection Machine
Principle Investigator: Lenore M R Mullin, Syracuse University

The development of Alpha-l, a basis for a parallel functional language is already underway. We have parallelized certain functions like the inner and outer product (based on the theory of 'the Mathematics of Arrays', by Lenore Mullin) and have run them on coarse-grained parallel architectures like the Encore's Multimax and Alliant's FX/80. Our next step is to extend the use of these parallelized operations to other computer architectures, to operating systems, and to compilers. We will test the efficiency and speedups of the scheduler for the parallelized array operations on fine-grained architectures like the Connection Machine and the hypercube. We intend to use the concepts of outer product and shared memory for comparison of process control blocks of an operating system. And we will study how to incorporate the parallelized optimizations of Alpha-l into a compiler.

Title: Parallel algorithms for discrete event simulation
Research area: Simulation of Large-Scale Systems
Architecture: C on the Encore Multimax
Principle Investigator: Robert G Sargent, Syracuse University

My research addresses two problems in the area of large-scale discrete event simulations: they typically have long computer run times, and they usually

are labor-intensive with respect to the development of the simulation models. One of my Ph.D. students and I are developing an approach for parallel discrete event simulation on a shared memory multiprocessor computer regarding these two problems. Our approach differs in three ways from most other research: 1) it is targeted for shared memory parallel computers; 2) it identifies when events are 'safe' to be executed in parallel; 3) we plan to evaluate it experimentally on an actual parallel computer. Our approach will include a parallel implementation of a small modelling language that we are developing for this purpose in C.

Title: Computational Electromagnetic Fields in Parallel Computers
Research area: Electromagnetic fields
Architecture: Fortran on the Alliant, the CM-2
Principle Investigator: Roger F Harrington, Syracuse University

The method of moments in electromagnetic field theory was developed by our research group in the nineteen sixties when high speed computers were finding many practical applications. Because of its simplicity, efficiency, and accuracy, this method has been widely used in electromagnetics engineering and has established itself as one of the standard methods in numerical analysis of electromagnetic fields. Since the matrix structure of the method of moments should allow us to develop parallel algorithms for this technique, we should be able to apply this technique to physical problems of increasing complexity, which could require hundreds or even thousands of expansion functions. Our current research is to consider the electromagnetic penetration into a conducting body of arbitrary shape through an arbitrary aperture. We wish to determine the aperture field distribution, the electric current in the conducting surface, the radar cross section, and the field inside the conducting body. We plan to calculate these quantities by modelling the conducting surface and the aperture with planar triangular patches and

computing the impedance of the aperture-closed conducting surface and other related matrices. The sources of parallelism in this technique are that the elements of the moment matrices can be calculated independently and other matrix operations.

Title: Mapping C and LISP into C* and *LISP
Research area: Parallel Languages, Automatic Parallelizing
Compilers
Architecture: Connection Machine
Principle Investigator: Mr. Shuo-Hsien Hsiao,

This project proposes to study the problem of building and interface to translate Lisp and C into *Lisp and C*, respectively, so that users may use parallel processing to solve their problems without understanding the architecture and network of processors on the Connection Machine. We propose to 1. Abstract the parallel elements of C* and *Lisp. 2. Find out the subset of C and Lisp languages that can be mapped easily into C* and *Lisp, respectively. 3. Design some rules to translate the more complicated C and Lisp structures into C* and *Lisp structures. 4. Construct a library for C and Lisp users to facilitate parallel computing.

Title: Nonlinear finite element study of the flow and
fracture of creep-damaging bars
Research area: Mechanical Engineering
Architecture: Alliant FX/80
Principle Investigator: Mr. Alan J Levy, Syracuse University

It is well known that the simple power law, which is often used to characterize the elevated temperature creep of polycrystalline metals, fails to predict the tendency for catastrophic neck growth which is observed in

constant load tensile tests prior to fracture. The problem as to what are the ultimate causes of necking in creep have still to be resolved, however, it appears likely that material damage in the form of grain boundary cavitation will significantly affect this phenomenon. It is the purpose of this project to study the growth of inhomogeneous deformation in materials in which grain boundary cavitation is significant. Since the problem is formulated with finite deformations with non-uniform stress and strain fields, the finite element method together with a finite deformation constitutive model are employed. The finite element program has been run on several computers; the fastest run time was 60 CPU hours on only a few elements. It is hoped that the accuracy of this numerical study can be improved with a computer with higher computational speed.

Title: Simulation, parallel programming and hardware modeling for data bases
Research area: Databases
Architecture: Connection Machine, Multimax
Principle Investigator: P Bruce Berra, Syracuse University

We have been performing several simulation, parallel programming and hardware modeling tasks on the Multimax and Connection Machines. We are still evaluating different schemes on those systems with specific applications to parallel computer architectures for data/knowledge base management. The Multimax is used for computationally intensive simulation of very large dynamic files. We are planning to use it for performance analysis of database machines and the development of parallel algorithms for relational operations. Our current work on the Connection Machine is the development of 'Parallel Surrogate Files Processing Algorithms' in the context of very large data/knowledge bases. A parallel relational algebra model is being developed on these machines for deductive

database systems. We are also investigating the possibility of using parallel architectures for the modeling of optical database machines.

Title: Computation of benchmarks for an optimally
efficient method of stochastic approximation
Research area: Vision
Architecture: Alliant
Principle Investigator: Denis Pelli, Syracuse University

Vision research shares with many other fields a heavy reliance on methods of stochastic approximation to estimate threshold, the critical intensity of a stimulus required for a standard level of performance by the observer. The problem is to choose the next test intensity in a series of tests so as to minimize the variance of the final threshold estimate after a given number of tests. I have formulated both an optimal solution, whose computational requirements increase exponentially, and a practical procedure which requires only a modest amount of computation. A supercomputer is required to run simulations of the optimal procedure to provide a benchmark for measurement of the absolute efficiency of the practical procedure, which I anticipate will be in excess of 90%.

Title: VLSI implementation of a semaphore memory
Research area: Computer Architecture
Architecture: Alliant
Principle Investigator: Sergio R R Chavez, Syracuse University

This project is to convert the wait and signal semaphore operations into simple memory access to variables of type natural number in order to simplify the implementation of critical sections in operating systems. The implementation will be achieved by designing a chip that will retain the pin

configuration and timing signals of normal memory chips but that will have internal features that will permit it to substitute the complicated wait and signal instructions by simple read instructions, leaving the task of updating the semaphore flags to the internal circuitry. In addition to designing (and fabricating) the chip, I intend to simulate operating systems problems such as the producer-consumer or the dining philosophers using the simplified critical sections and describe and test an application of the chip in the implementation of a process scheduler in a coarse grained parallel computer.

Title: Parallel sorting with serial comparators
Research area: Parallel Algorithms
Architecture: *Lisp and Paris on the Connection Machine
Principle Investigator: R Vaidyanathan, Syracuse University

Traditionally, parallel sorting of n m -bit keys consider m -bit comparisons of keys as the basic unit of time. Thus, if we look at the input to a parallel sorting algorithm as an $n \times m$ array of bits (representing n m -bit numbers), traditional approaches operate on entire rows of the input at a time. This does not take advantage of the fine-grained nature of the input. We propose an approach which processes the input, column by column. The advantage of this approach is that the sorting time is not constrained by the size of the comparator in the processors. Two parallel algorithms based on this approach have been implemented on the Connection Machine. The first uses n processors to sort n m -bit numbers in $O(m \log n)$ time. The second algorithm requires n^2 processors and uses $O(\log m + \log n)$ time.

Title: Parallel computation of EM fields to be applied in
high speed integrated circuit design
Research area: Electromagnetic Field Computation
Architecture: Alliant
Principle Investigator: Madhavan Swaminathan, Syracuse University

Electromagnetic scattering from surfaces of arbitrary shape involve the solution of complex integro-differential equations. The method of moments provides us with a technique to convert these equations into a matrix equation, which can be solved for the unknowns. This technique uses a two step process: 1. Generation of the complex matrix using a numerical procedure. 2. Solution of this matrix involving 1000's of unknowns. Parallel computing can reduce the time needed to generate and solve the matrix.

Title: A parallel language interpreter and garbage collector
Research area: Parallel language implementation
Architecture: C on the Encore, C* on the CM
Principle Investigator: Ms. Xue Li, Syracuse University

One of the more promising sources of parallelism in programming language implementation is in the implementation of functional languages. In this project, an implementation of a parallel functional language interpreter that can also run in parallel with one or more garbage collectors will be explored. The interpreter that will be implemented will be a simple lambda-calculus graph-reduction interpreter as described by Simon Peyton-Jones. The garbage collector will be a variation of the on-the-fly garbage collector originally proposed by Dijkstra. This kind of garbage collector can run in parallel with the interpreter using shared memory without any need for explicit locking. The synchronization is all hand'ed by the marking of the garbage collector. Fine-grained parallel graph-reduction interpreters, using Hudak's techniques, and garbage collectors may also be explored.

Title: Advanced Computing for Molecular Studies and
Design of Materials
Research area: Chemical Engineering
Architecture: CM, Alliant, Encore
Principle Investigator: Cynthia S Hirtzel, Syracuse University

Design of advanced materials, processing of rheologically and structurally complex fluids and similar operations of fundamental scientific and technological importance require sophisticated theoretical and experimental tools to probe molecular phenomena in such systems. For example, it is important that the role of molecular and supramolecular interactions on the formation and behavior of various phases and structures within high-strength ceramics, glasses, and electronic and optoelectronic materials be studied systematically. It is proposed in this project to develop Monte Carlo and Brownian Dynamics computer experiment packages; to develop algorithms for simulating configurations, motions and structural features and transitions in advanced materials; and to develop numerical packages for simulating the time-evolution of systems governed by stochastic differential equations.

Title: Implement & test address block location system on
Connection Machine including image analysis tools
Research area: Real-time image analysis
Architecture: C* and *Lisp on the Connection Machine
Principle Investigator: Mr Sargur N Srihari, SUNY Buffalo

Recently, software for a new generation of Address Block Location Systems (ABLS) has been designed to run on standard sequential computers. An ABLS locates and reads the destination address region on postal images

including letters, flats (e.g. magazines), and irregular parcels and pieces. The SUNY Buffalo ABLIS is based on a computational theory with a spatial model and an attribute model for postal images and uses many image analysis tools such as grey level adaptive thresholding, bottom-up segmenter for machine-generated text, regularity discrimination, handwriting analysis, and shape analyzer, among others. This system reaches an acceptable location rate of 81% on all mail streams (present OCR machines have a rate of 55% for letters only), but some of the components of the system take up to 30 minutes of running time. We plan to use a fine-grained parallel machine to investigate reducing the running time of the system by several orders of magnitude.

Title: Expert System for the Synthesis and Verification of
Systolic Architectures
Research area: VLSI design
Architecture: Encore Multimax and the Connection Machine
Principle Investigator: Camille Batarek, Syracuse University

VLSI technology allows many algorithms in pattern recognition and image and signal processing to be implemented on array processors. These array processors, called systolic arrays, consist of connected processing elements and use pipelining and multiprocessing to achieve high data rates. Several methodologies have been proposed which synthesize systolic arrays by transforming an algorithmic specification into a hardware design, but which cannot be fully automated because they rely on some steps that require reasoning by the designer. It is proposed here to use expert systems to implement the human reasoning that would fully automate the synthesis of systolic arrays from algorithmic specification, as well as to verify the correctness of an already available architecture.

Title: Image processing by simulated annealing
multiprocessor and parallel processors
Research area: Image Processing
Architecture: Connection Machine and Encore Multimax
Principle Investigator: Mohamed Zerrouki,

Simulated annealing is a newly developed technique to solve very complex optimization problems. It can find the global optimum (or near global optimum) solutions to a constraint criteria function with multi-local optima. It originated from the analogy between annealing of solids and the optimization of large combinatorial problems. It is proposed to use simulated annealing for the restoration of degraded images. The necessary components for this are configuration space description, a criteria function, a perturbation mechanism, and a cooling schedule. These techniques have been studied on serial computers. It is proposed that they be studied in the context of parallel computation.

Title: Simulation - spanning multiaccess channel
hypercube
Research area: Performance analysis of parallel architecture
Architecture: Encore Multimax
Principle Investigator: Patrick Dowd, Syracuse University

The overall goal of this simulation effort is to provide accurate estimation of performance characteristics of the Spanning Multiaccess Channel Hypercube (SMCH). The simulation will be developed in four phases, where at each phase, software emulation of the hardware is written in such a way as to isolate hardware components in order to minimize any impact on the existing code with a design change. Ultimately, it will simulate a highly parallel computer architecture with an extremely large number of nodes. This approach makes it possible to study variations in performance

characteristics with variations in target hardware. This enables us to study the underlying causes of performance degradation.

Title: Parallel compilation using pipeline and data partitioning
Research area: Compiler Design
Architecture: C on Encore Multimax, C* on Connection Machine
Principle Investigator: Sanjay Khanna, Syracuse University

We propose to develop a parallel compiler based on a new data partitioning scheme. The scheme is based on partitioning the grammar of the language being compiled, which in turn effectively decomposes the language specification into multiple subsets. The scheme requires definition and development of specialized subcompilers. These would be implemented by a simple and practical method using Parser Generators. A comparison with other methods such as functional decomposition and previously proposed data partitioning methods, reveal that the proposed methodology has a potential speed-up factor larger than that obtained by functional decomposition. The overall structure of the compiler would also achieve parallelism by pipelining the three stages: lexical analysis, syntax analysis and code generation. We plan to use these compiler design ideas to implement a parallel compiler for Pascal on the Multimax and C on the Connection Machine.

Title: Parallel Algorithm for P. version finite element analysis
Research area: Stress Analysis (numerical analysis)
Architecture: Fortran on the Alliant FX/80
Principle Investigator: Eric M Lui, Syracuse University

The emergence of parallel multiprocessor computers has revolutionized the design and implementation of computer algorithms. One of the most commonly used numerical methods for stress analysis is the finite element method. In this method, the structure or structural component is discretized into elements of finite size, interconnected at nodes. Although this procedure is theoretically sound and well-developed, the implementation of this algorithm for large-order problems is often prohibitive with conventional serial computers. In view of this, an efficient algorithm referred to as the parallel conjugate gradient method is proposed. This method is very adaptive to a parallel environment because of the inherent parallelism of the approach. Each element or group of elements is processed by a separate processor, with communication between processors occurring only intermittently.

Title:	Image Processing, pattern recognition of features in the image
Research area:	Image Processing
Architecture:	C on the CM1, CM2, and the Multimax
Principle Investigator:	Elias Mazur, Syracuse University

The proposed research concerns the processing of large satellite images. The processing includes enhancing and extracting features of such images to be further recognized and classified. The goal is to be able to extract and recognize storms and also predict storm movements. A comparison will be made between the relative performance of the Connection Machine and the Encore Multimax when applied to this kind of image processing.

Title: Implementation of Head-order Reductions of
Lambda-Expressions on Connection Machine
Research area: Logic Programming
Architecture: Connection Machine
Principle Investigator: Sining Zhang, Syracuse University

Work I have done so far: 1) The head-order reduction scheme generates head-order normal forms from initial lambda-expressions, and recursively all arguments are then concurrently reduced to head-order normal forms. The following parts of the reduction have been implemented: unwind-to-head, eta-extension-in-the-large, beta-reduction-in-the-large, identity-reduction-in-the-large. 2) The environment is simply attached to and shared by the expressions, and methods of building and cutting back the environment are implemented. I propose to do the following: modify the reduction program developed so far, optimize the global control and communication, and implement sharing of the environment and the previous reduction. entation of parallel head-order reduction of lambda-expressions.

Title: Maintain macros for parallel programming
developed at Argonne for Encore and Alliant.
Research area: Parallel Programming
Architecture: Alliant FX/80 and Encore Multimax
Principle Investigator: Nicholas Karonis, Syracuse University

As a summer programmer at Argonne National Labs, Karonis debugged and maintained the set of parallel macros they had written for their Encore, Sequent, Alliant, and Suns. These macros allow user to write parallel programs that consist of C code with macro calls for the parallel constructs. A version of the macros was written for each machine so a C program with embedded macros is perfectly portable to each of these machines. These

macros will be installed and maintained on the Alliant and the Multimax at NPAC.

Title: Implementation of the Image Algebra on the
Connection Machine
Research area: Machine Vision
Architecture: *Lisp on the Connection Machine
Principle Investigator: Joseph N Wilson, University of Florida

The Image Algebra Project, funded by AFATL/Eglin, resulted in the selection of a design -- developed by the University of Florida -- as AFATL's Image Algebra. This algebra, which is both concise and mathematically complete, is a mathematical formalism for describing image processing transformations. The Image Algebra can be executed very rapidly on specialized parallel architectures, but a general interconnect, massively parallel machine like the Connection Machine is an ideal machine architecture on which to execute algorithms described in the algebra. This is especially true for the Image Algebra's generalized template operations -- extremely flexible neighborhood-like operations giving great expressive power. We propose to implement the Image Algebra operations in their full generality on the CM-2. Once implemented, these operations should make it easier for novice CM users to develop computer vision algorithms to run on the CM.

Title: Parallel Algorithms for NLSI Testability Analysis.
Research area: VLSI Design
Architecture: Encore Multimax
Principle Investigator: Dr Gary L Craig, Syracuse University

One of the testing challenges of static CMOS logic networks is the detection of stuck-open (s-op) faults which cause combinational circuits to exhibit

sequential behavior. Testing schemes for these faults require multiple-pattern tests. In an attempt to solve the problems of previous test generation algorithms, a minimum length robust multiple-pattern test generation/testability analysis procedure has been developed. This process involves creating a database for the network of representative cubes for each circuit node and then generating a test pattern for every 'fault' in the network. Both of these procedures are well suited for parallel computation.

Title: SUMAC Project
Research area: Computer architectures and applications
Architecture: Alliant, Connection Machine, Multimax, Sun Lab
Principle Investigator: Charles D Stormon, Syracuse University

The SUMAC (Syracuse University Machines for Associative Computation) Project at the Syracuse University CASE Center is a long-term, broad-scope venture into fine-grained parallel processing computer architectures and applications. We are currently developing applications and implementations of associative architectures. Some implementations use several prototypes of our custom VLSI Content Addressable Memory (CAM) chip, while others are simulated using the Alliant and the Connection Machine. Current applications work is centered on Logic Programming, Image Processing, Expert Systems for Automated Test, and Knowledge-base Management Systems.

Title: Parallel algorithm of Quicksorting
Research area: Parallel Algorithms
Architecture: C on the Multimax, C* on the Connection Machine
Principle Investigator: Fuxu Cao,

This project is to evaluate two different parallel implementations of a traditionally sequential algorithm, QUICKSORT. The implementations will be done on computers of SIMD and MIMD architecture and measurements will be taken of their performance.

Title: An Adaptive learning model of Organizational Information Systems.
Research area: Social and Decision Sciences
Architecture: *Lisp on the Connection Machine
Principle Investigator: Serge Taylor, Carnegie-Mellon University

It is a difficult problem to predict the consequences of changes in organizational information and control systems since the formal tools available for studying organizations are not well suited to analyzing the mechanisms by which information produces its effects. The major innovation of this project is to conceptualize changes in organizational information systems as a process involving adaptive learning based on organization members' experiences. The model for the learning is an artificial intelligence methodology called 'learning classifier systems' with some changes to fit the organizational model. Individuals are represented as having beliefs that change as a result of experience, with those hypotheses which experience shows are more productive becoming more likely to be used in the future. Beliefs are represented as decision rules, and entirely new decision rules can be formed from pieces of old ones as a result of experience. This project proposes to carry out systematic testing of different organizational representations and of modifications to the basic classifier system algorithms.

Title: Implementing high precision arithmetic on the Connection Machine.
Research area: Computer Arithmetic
Architecture: *Lisp on the Connection Machine
Principle Investigator: Ms. Mobolanle Odulate, Syracuse University

The objective is to implement high-precision arithmetic on the Connection Machine, making it possible to manipulate numbers which are too large for any one processor. The numbers' size will be limited only by the total number of processors available. Specifically, parallel algorithms for the addition operation will be implemented, and experiments will be conducted with the timing of the various algorithms.

Title: Parallel Algorithms for image processing and image synthesis/parallel algorithm for the network flow problem
Research area: Computer Vision
Architecture: C* on the CM, C on the Multimax, C on the FX/80
Principle Investigator: George A Betzos, Syracuse University

In this project, parallel algorithms for image processing will be developed. First of all, fine-grained parallel algorithms for the two dimensional Fast Fourier Transform, image enhancement algorithms (smoothing, sharpening) and segmentation algorithms (relaxation) will be developed for the Connection Machine. Furthermore, the combination of ray tracing and the radiosity approach for the generation of highly realistic images will be implemented using MIMD techniques.

Title: Parallel Algorithms for VLSI layout compaction.
Research area: VLSI Design
Principle Investigator: Sachin Sapatnekar, Syracuse University

The idea of VLSI is to pack in as much as possible into as little space as possible to obtain the best possible space efficiency. At the same time the physical layouts of VLSI systems must obey certain design rules that dictate physical constraints in the spacing between adjacent wires or devices. With the advent of more and more complex circuits, the need for computer aided design tools that handle problems that were previously handled by human designers has been underlined. Parallel architectures play an important role in computer aided design of such tools. The basic approach to parallel layout compaction is to divide the layout region into a number of parts that may be allotted to different processors. Each of these sub-regions is separately compacted using conventional serial compaction algorithms. Next, these sub-regions must be integrated into a whole that satisfies the boundary conditions that prevails between them. This is an iterative process that must be repeated several times before a final result is obtained.

Title: Parallel Constraint Satisfaction
Research area: Machine Vision
Architecture: C* on the Connection Machine
Principle Investigator: Michael J Swain, University of Rochester

Many problems in Artificial Intelligence can be posed as constraint satisfaction problems (CSP's), in which the possible labels for a set of variables must be determined subject to a set of constraints. One of these problems is the graph isomorphism problem, which we are interested in for object recognition. Sequential algorithms have been developed and analyzed for arc consistency, a method for eliminating local inconsistencies that cannot possibly take part in the global solution to the constraint satisfaction problem.

We have developed a parallel version of this algorithm, a SIMD message-passing algorithm for arc consistency that obtains linear worst case speed-up over the best sequential algorithm. We are implementing this on the CM-2 and plan to test its performance on large constraint satisfaction problems.

Title: Application of Artificial Neural Networks for
Learning Control of Robot Motion.
Research area: Robotics
Architecture: Connection Machine
Principle Investigator: Mr. M. K Ciliz, Syracuse University

Autonomous control of robots requires the utilization of sensory information at various levels of complexity that are, in general, descriptive, rather than purely numeric. Hierarchical controllers have been proposed to solve the path planning, navigation, obstacle avoidance, and motion control problems. Rule-based methods have been predominantly utilized since they allow symbolic processing of descriptive information. Since the correctness of the rules depends heavily on the environment, changes in the environment should produce changes in the control rules. In this project, the utilization of neural net architectures for adaptive-learning control is considered. The conceptual design of a two-level learning system will be investigated, and the developed system will be implemented on a parallel machine.

Title: Parallel Computing Application Study (SCEE &
AFWL/SC, Kirkland AFB) [as CASE members]
Research area: Numerical Analysis
Architecture: C, Assembler, etc. on the CM2 and the FX/80
Principle Investigator: Fred H Schlereth, Syracuse University

The purpose of this research is to experiment with efficient methods for numerical analysis on the FX/80 and the Connection Machine (CM2). Currently, algorithms and architectures for the solution of very large matrices (10,000 by 10,000) are being investigated. Data flow between the processing elements and memory is the main problem in taking advantage of parallelism for scientific computing problems, especially as the size of the problems gets large. One aim of the project is to apply computing strategies developed for ASIC parallel computers to these matrix computations.

Title: Implement a fast parallel graph reduction on
Connection Machine

Research area: Logic Programming

Architecture: *Lisp on the CM

Principle Investigator: Feng Yang, Syracuse University

I would like to do the following: 1) develop an algorithm to transform lambda-expressions into SKI-combinator expressions for an efficient implementation of graph reduction; 2) investigate the possibility of parsing lambda-expressions in parallel; 3) develop a graph reduction scheme to implement an extension of SKI-combinators in parallel; 4) develop a combination of reference-counter and mark-collecting garbage collection methods and implement it on the Connection Machine.

Title: Implementing and testing various simulation
models of neural network computing.

Research area: Neural Nets

Architecture: Connection Machine

Principle Investigator: Ming Li, Syracuse University

Massively Parallel distributed processing has been studied to deal with various discrete dynamical system problems, either theoretical or practical. However, there is no proper theory to predict the dynamical behavior of this sort of system. Most of the research works have to resort to simulation, even for some purely theoretical topics. In order to explore a theory of behavior, however, it will be necessary to develop a work-bench or test-bed for problems of this sort. It is the goal of this research to start developing these tools by programming several neural-net models on the CM.

Title: Relational databases on a cube-connected architecture
Research area: Relational Databases
Architecture: C on the Encore Multimax.
Principle Investigator: Mr. Paul Jackson,

In this project, several innovative parallel algorithms will be implemented and analyzed which compute relational database operations such as select, project, and join. These algorithms were developed for a boolean cube-connected architecture which will be simulated on the Encore Multimax. Each of these operations tries to alleviate the problem of non-uniform data distribution amongst the nodes of the cube by redistributing the data as part of the algorithm.

Title: Parallel channel routing algorithm for SIMD processor.
Research area: Computer Aided Design
Architecture: Connection Machine
Principle Investigator: Shantanu Ganguly, Syracuse University

Routing problems in the area of Computer Aided Design consist of a set of modules along with a list of their interconnections, which have to be

accomplished while satisfying certain constraints. Channel routing is a special case of routing where terminals are provided only on the top and bottom of a square with a list of interconnections between terminals. Several sequential algorithms exist for channel routing, which are very simple and fast, while the quality of the solution produced by this method is not very good. Unfortunately, sequential algorithms that do produce good solutions have not proven amenable to parallelization. Our approach is to use an algorithm of the first type to obtain an initial solution that is not very good. The initial layout can then be compacted in area using fine-grained parallelism, as well as post processing such as removing unnecessary wires and bends, to achieve a final high quality solution.

Title: Scheduling of Distributed Computing Systems.
Research area: Computer Operating Systems
Architecture: C on the Multimax and the CM2
Principle Investigator: Ishfaq Ahmad,

The proposed research is related to the design of task-scheduling algorithms for distributed computing systems. We intend to analyze different dynamic scheduling algorithms for distributed homogeneous computer networks which consist of large number of independent computing nodes connected with each other according to a certain topology. The performance of some existing algorithms will be compared with the newly proposed schemes. The comparison will be done by extensively simulating the algorithms and collecting data for various performance indices such as system's throughput, response time, etc.

Title: Parallel Implementation of Accelerated Algorithms
for Monte Carlo in Statistical Mechanics and
Quantum Field Theory.
Research area: Theoretical High Energy Physics/Statistical
Mechanics
Architecture: CM2
Principle Investigator: Mr. R. Brower, Boston University

Monte Carlo simulations have become an increasingly important tool in studying complex statistical systems and fundamental properties of quantum field theory, as well as providing a pseudo-dynamics to probe quasi-equilibrium. A particularly important problem with the dynamics of the Monte Carlo algorithm is that all local algorithms suffer from 'critical slowing down' as one approaches a phase transition. To avoid critical slowing down, one must define non-local (or collective) variables and a new dynamics for driving them. Such non-local algorithms have proven difficult to vectorize. They appear to be natural candidates for implementation on fine-grained parallel machines such as the Connection Machine. We are using these new Monte Carlo techniques in the study of critical properties of magnetic systems and of non-Abelian gauge theories for particle physics.

Title: Parallel Molecular Dynamics Simulations
Research area: Molecular Dynamics
Architecture: CM2
Principle Investigator: Mr. R. Giles, Boston University

This proposal is to develop and exploit new algorithms for doing parallel molecular dynamics simulation of dense fluids. The main interest is to understand the microscopic processes involved in the evolution of supercooled liquids to either crystalline solids or glasses. Since the early stages of this evolution cannot be detected in laboratory experiments,

computer simulations of this evolution are of much importance. At present, the primary focus of this research is to understand the nature of nucleation for deep quenches where the classical description of nucleation is not applicable. Some molecular dynamics simulations of a simple Lennard-Jones liquid that have already been done. Molecular dynamics simulations require the integration of Newton's equations for a system of particles. The problem here differs significantly from that of bonded molecular interactions in the the relevant neighbors of a given particle change dynamically as the system evolves and cannot be determined statically at the beginning.

Title: Parallel Micromagnetic Simulations
Research area: Electromagnetics
Architecture: CM2
Principle Investigator: Mr. R. Giles, Boston University

The dynamic properties of magnetic domains in thin films are of great interest both theoretically as an example of nucleation phenomena and practically for applications to magnetic recording. This research is to implement dynamic micromagnetic simulations of vertical Bloch line structures in multilayer films. These materials can be described classically by a field of 3 component dipole moments whose time evolution is described by the phenomenological Landau Lifshitz Gilbert equation. They are approximated for simulation by a regular rectangular or hexagonal lattice structure. One time step in the simulation is effected by computing the local magnetic field, as a function of the local dipole moment, the nearest neighbor moments, and of the long range magnetostatic field, at each point in the lattice. From this is calculated the velocity of the dipole moment, which is integrated using standard methods. Typical studies of a single sample require 2000 to 20000 time steps.

Title: Simulation of Dynamical Behavior of Model
Generation Operators in the MGR Architecture
Research area: Computational Logic
Architecture: *Lisp on CM1, CM2
Principle Investigator: Dr. M. J Coombs, New Mexico State University

The purpose of this project is to identify dynamic control strategies for the Model Generative Reasoning (MGR) problem solving architecture. The use of continuous-valued dynamical systems to specify control, instead of the symbolic heuristics or algorithms used in traditional knowledge-based systems, has the potential to greatly increase the flexibility of the behavior of automated problem solvers in response to the data and knowledge limitations imposed by unstructured task environments. MGR consists logically of six operators, each of which is a SIMD machine, organized in a shared memory parallel configuration. A dynamical system for controlling MGR by varying the rates of execution of the MGR operators, which is defined over variables that measure the syntactic properties of MGR knowledge structures, has been defined. A simulation environment with which to investigate the behavior of this dynamical system in different data and knowledge environments experimentally is currently being implemented in *Lisp on the Connection Machine.

Title: Two-Dimensional Kalman Filtering Using Parallel
Processors
Research area: Image Processing
Architecture: FX/80, Multimax, and Connection Machine.
Principle Investigator: Mr. David J Potter Professor, Clarkson University

This project involves the implementation of two-dimensional Kalman filtering on different parallel architectures. Currently, Sequent and CLIP machine implementations have been studied, but it would be useful to study

implementations on the architectures available at NPAC. Two-dimensional Kalman filtering is computationally intensive. However, unlike most image processing algorithms, the parallelism is not simply on the data at each pixel which would give machines such as CLIP (processor arrays) the best advantage for optimization). The Kalman filter is based on the scan-line model such that a pixel in the middle of the image is a function of pixels above and to the left of that pixel. For the Encore and Alliant machines, we feel the changes in the algorithms used on the Sequent will be slight. For the Connection Machine, new algorithms will have to be devised and implemented.

Title: SIMD Massively Parallel Exhaustive Search of Small State Spaces
Research area: Artificial Intelligence
Architecture: C/Paris on the CM-2
Principle Investigator: Mr. Lewis Stiller, Boston University

Exhaustive state space search can yield optimizations which are qualitatively different than heuristic search would be likely to yield. This technique has already been used to find new solutions to chess end games with 5 pieces and to optimize assembly language programs. This research proposal is to implement an exhaustive search of a game space by computing all possible moves from a game configuration. This method achieves greater functional simplicity at a cost of considerable computational redundancy. The CM-2 will be used for the implementation, which will permit heretofore infeasible real-time domain interaction and experimentation. This may help use endgames to explore induction learning and expert systems. This project will also explore the canonical exponentially time complete game Peek, a game played on propositional formulae.

Title: Development of Parallel Algorithms for Position
Control of Cooperating Robot Arms.
Research area: Robotic Engineering
Architecture: Fortran on the Alliant
Principle Investigator: Manolis Christodoulou, Syracuse University

This research is motivated by problems encountered in position control of cooperating robot arms. The objectives of this work are the development of efficient path tracking algorithms and the application of parallel processing techniques. The proposed algorithms are direct, i.e. do not require the inverse kinematics, guide the manipulator in arbitrary paths, and satisfy proximity criteria. The use of many processing elements will increase speedup and efficiency. The viability of the approach will be demonstrated by verifying the algorithm experimentally. Existing laboratory robots will be connected on line with the parallel machines for this purpose.

Title: Performance evaluation and modeling of
multiprocessor architectures
Research area: Computer Architecture
Architecture: C* on the Connection Machine
Principle Investigator: Daniel J Pease, Syracuse University

The design of a machine architecture would be considerably aided by being able to simulate the running of the machine under a variety of assumptions about the machine's use and the algorithms used in the machine. By using the isomorphism between Petri nets and dataflow graphs, we approach this problem by having the architect describe the machine using dataflow graphs and actually running their architectural designs on the Petri nets. Timed Petri nets have already been implemented on the Connection Machine, and preliminary results show that machine designs with 1000 processors and 2000 memory modules can be run in less than an hour. (Previous net

implementations on sequential machines ran 16 processor designs in 9 hours.) We have also validated the Petri net implementation by modeling some existing architectures; namely, the Connection Machine, Alliant, and the Multimax. The results are consistent with the actual running of the machines.

Title: Proofs of correctness for *Lisp
Research area: Parallel programming languages
Architecture: Connection Machine
Principle Investigator: Dr. Nancy J McCracken, Syracuse University

One of the important problems of parallel computing is to design programming languages that allow programmers to specify their problem with ease and reliability and also have an efficient implementation on a parallel machine. We propose to study a parallel programming language that has been designed and implemented on a massively parallel computer, the Connection Machine, to see what features of a data parallel language meet the requirements of good software engineering, and what does not. Our method of study has been to work on an inference system for proofs of correctness for *Lisp. Writing an inference rule for each statement in the language specifies the semantics of the language very precisely and shows the semantic ideas that a programmer must deal with. Furthermore, using the rules to write proofs of correctness for programs written in *Lisp, should show what notations and concepts are important for programmers to write good specifications of problems that are suitable for data parallelism.

Title: Tools and techniques for Connection Machine programming
Research area: Parallel algorithms
Architecture: *Lisp and PARIS on the Connection Machine
Principle Investigator: Mr. Bill O'Farrell, Syracuse University

The goal of this project is to investigate basic methods and tools for programming the Connection Machine. We have developed new techniques for processor allocation and reclamation which are essential for the implementation of various programming languages on the Connection Machine. In addition, we have found ways to optimally exploit the unique features of the connection machine to speed the execution of a problem in number theory known as the 'Syracuse' (or $3x + 1$) problem. We will also use the proof techniques developed by N.J. McCracken to prove the correctness of the algorithms.

Title: The Syracuse Problem
Research area: Mathematical sequences
Architecture: *Lisp on the Connection Machine
Principle Investigator: Mr. Ken A McVearry, Syracuse University

The Syracuse problem, also known as the $3x + 1$ problem, involves a conjecture about a particular mathematical sequence that has yet to be proven. The sequence is defined by the following two rules: if the starting value is odd, multiply it by 3 and add 1 to find the next value in the sequence; if it is even, divide it by 2 to find the next value. For example: 1, 4, 2, 1. Every starting value (positive integer) that has been tested converges to 1 as in this example, yet the theorem that this is true for all starting values has yet to be proven in spite of intensive activity in this area. We would like to use the Connection Machine to find a counter-example to this theorem, since it has resisted proof by reasoning. We have already accomplished our secondary

goal, which was to break the known world record for the highest starting value tested.

Title: Extended precision arithmetic on the Connection Machine
Research area: Semi-numerical algorithms
Architecture: *Lisp on the Connection Machine
Principle Investigator: Mr. Ken A McVearry, Syracuse University

Some basic integer arithmetic operations on the Connection Machine are limited to operands of relatively small size, e.g. 128 bits. There are many interesting projects that could be explored using the CM if arithmetic could be performed on much larger numbers. To this end, we would like to write *Lisp functions that would perform operations on very large numbers stored across many processors. This will involve the discovery and implementation of new algorithms for performing arithmetic on data that is spread across processors on a data-parallel computer. These algorithms should be $O(\log n)$, where n is the number of processors used to store numbers. Operations to be implemented include addition, subtraction, and multiplication.

Title: QCD and 3-d Ising Project on CM-2 and CM-1
Research area: Theoretical Physics
Architecture: *Lisp on the Connection Machines
Principle Investigator: Clive Baillie, California Institute of Technology

We are using the Connection Machines at NPAC to investigate two lattice models. On the CM-2 we will run a lattice gauge theory simulation of quantum chromodynamics (QCD), and on the CM-1 we will investigate the critical properties of the 3d Ising model using the Monte Carlo Renormalization Group (MCRG) method. Simulation of QCD are one of the

most floating-point operation (flop) intensive numerical calculations performed by physicists. Hence, the CM-2 with its large number of Weitek floating-point accelerator chips (1024 for the 32K processor machine at NPAC) is an ideal machine to use for this calculation. On the other hand, the Ising model is the simplest statistical-mechanical spin-model making use of single bit (i.e. logical or boolean) variables. So it is a problem well suited for the CM-1's bit-serial processors (32K at NPAC).

Title:	Calculations of the 3-D structures of proteins
Research area:	Molecular Bioscience
Architecture:	C* on Connection Machine, Fortran on Alliant and Multimax
Principle Investigator:	Harold Scheraga, Cornell University

The goal of this project is to predict the three-dimensional structure of relatively small proteins (50-60 amino acid residue range) by using algorithms to locate the minimum of potential energy that corresponds to their native conformations. Based on the hypothesis that the native conformation of a protein is determined completely by its constituent sequence of amino acids, statistical mechanical arguments show that the native conformation must correspond to a minimum in the potential energy surface that governs the motion of the protein in its physiological environment. Computational models will be generated which should closely predict the most physiologically stable and therefore biologically consistent protein structure. The first of two algorithms to be used is based on building an amino acid chain by combining small pieces for which all minima are located within a fixed energy of the global minima. These are in turn combined on the basis of locating the low-energy minima of the new fragment, and so on in tandem. The second algorithm uses an adaptive Monte Carlo method, the partition function, and other integrals of statistical mechanics to obtain the average and most probable conformations directly.

Title: Distributed computation of Image Labelings
Associated with Markov Random Fields
Research area: Machine Vision
Architecture: C* on the CM-2
Principle Investigator: Michael J Swain, University of Rochester

Markov Random Fields (MRF's) have been successfully applied to image and surface reconstruction and segmentation problems in Computer Vision. A major problem with MRF's is the computational intractability of calculating the exact solution to the problem they define. Among the algorithms which have been developed to calculate a close to optimal solution is the Highest Confidence First (HCF) algorithm, which is a deterministic sequential algorithm that chooses labels in order from highest to lowest confidence. We hope to develop a parallel version of this algorithm by exploiting two sources of parallelism: the relative confidence relation is a partial order and the siblings in the resulting directed acyclic graph can be computed in parallel, and a parallel heuristic choice of 'good' labels may be possible. We plan to implement the resulting algorithms on the Connection Machine and compare its results and timings with a variety of algorithms for solving MRF's.

Title: An interactive information retrieval system based on
a connectionist model
Research area: Information Retrieval
Architecture: C* on the Connection Machine
Principle Investigator: Robert N Oddy, Syracuse University

The objective of this work is to explore a type of automatic information retrieval system which interacts with its user by building, and continuously

adjusting, a stored image of his/her information need. This contrasts with the usual approach of matching a clearly expressed query with portions of structured records. We are resuming earlier work on a prototype information retrieval system, called Thomas, which was able to find documents relevant to the user on the basis of an extremely sketchy initial query and evaluative reactions given by the user as the dialogue proceeded. Thomas' processes are intrinsically highly parallel, akin to spreading activation in large networks. The C* programs we are developing manipulate networks of words, names, and documents. The texts are handled in conventional ways, in indexed files on the VAX host computer, and the Connection Machine is used for the operations on the network structure.

Title: Transmission lines.

Principle Investigator: Mr. Tawfik R Arabi, Syracuse University

The objective of this research is to develop algorithms to fully characterize multiconductor transmission line networks. The number of networks (possibly non linear) and transmission lines connecting them are arbitrary. The lines are lossy and embedded in multilayered, homogeneous, lossy dielectric media but have uniform cross sections along their lengths. In our formulation, we utilize the QUASI-TEM assumption, therefore the lines behavior can be described in terms of their frequency dependent circuit parameters $[R(f)]$, $[L(f)]$, $[C(f)]$ and $[G(f)]$.

Title: Geological remote sensing & image processing in parallel.

Research area: Geology

Architecture: *Lisp on the Connection Machine

Principle Investigator: Bronya Oldfield, Syracuse University

The research proposed in this project is to determine if subsurface oil reserves can be located in satellite images. The technique involves analyzing filtered, enhanced images of multispectral satellite data to find trapped hydrocarbons. The image processing is computed on the Connection Machine and includes algorithms for Fourier and inverse Fourier transforms. Using the CM reduced processing time for a 512×512 pixel image of single band spectral data from one hour on a serial computer to five minutes. Initially, the technique has been tested on an database containing cultural features, oil and gas wells, geological and topographical data of southwestern Wyoming, supplied by ARCO Oil and Gas Company. The initial results show known oil fields clearly. The next step is to verify the image processing techniques in unexplored areas to determine locations of previously unidentified oil and gas reserves.

Title:	Modifications to the APL2 language to support expression and execution of generalized algorithms
Research area:	Parallel Languages
Architecture:	CM, Multimax, Alliant
Principle Investigator:	Mr. Robert G Willhoft, IBM Corporation

The APL language has received attention in the last several years as a language that can be used to express parallel algorithms. The primary focus of interest has been its ability to express algorithms on vector or array arguments directly without the need for the programmer to convert them into sequential loops as in most languages. Despite this very significant processing advantage, both APL and APL2 do not handle a number of other significant parallel structures that are necessary for efficient, general parallel behavior. APL tends to be very good at expressing fine grain parallelism, but has few facilities to express course grain parallelism or data flow structures. This research proposal focuses on some of the enhancements that would be necessary to the APL2 language to provide some of these missing features,

and intends to implement critical subsections of the new language on various parallel architectures as a feasibility study.

Title: Extracting program structure knowledge from
Fortran programs for use in converting to parallel
algorithms

Research area: Parallel Languages

Architecture: CM, Alliant, Multimax

Principle Investigator: Mr. Robert G Willhoft, IBM Corporation

There has been a great deal of research done on the automatic conversion of FORTRAN programs to run on parallel machines. This work includes a broad spectrum of work ranging from vectorizing compilers to the extraction of complete data flow algorithms from sequential algorithms. Much of this work focuses on the conversion of FORTRAN to a specific machine. However, most of these methods concentrate on extracting static data from the program. In doing this much of the possible parallelization is lost because of data dependencies that are unknown to the static program analyzer. This means that in any case the worst possible assumptions must be made. On the other hand, dynamic analysis of a program is time consuming and subject to misleading information due to the problem of selecting the correct input data. It is the opinion of this researcher that by extracting knowledge from the FORTRAN code it should be possible to gain most of the advantages of dynamic analysis without actually having to run the code. This proposal deals with specific Research areas that should be looked at to extract the necessary knowledge.

Title: Parallel Processing in Knowledge-based Computer Vision Systems
Research area: Machine Vision
Architecture: Connection Machine, Alliant FX/80, Encore Multimax
Principle Investigator: Dr. Minsoo Suk, Syracuse University

We are concerned with the development of robust techniques for 3-D object recognition procedures dealing with multiple-object scenes with partial occlusion and background clutter. Our problem at present is limited to rigid objects with six degrees of freedom. We would extend our results to rigid objects with parameterized variations, rigid objects in motion and deformable objects. The ultimate goal of our research is the better understanding of the following theoretical and practical issues: development of suitable indexing schemes for indexing into a library of object models represented at multiple levels of abstraction, insight into how symbolic and numerical constraint satisfaction techniques could be integrated into a problem-solving mechanism, and a more robust and general framework for 3-D object recognition.

Title: Manipulation & implementation of the Hamming Net for Pattern Classification, Signal Processing—using the elements of Prediction and System Modelling
Research area: Signal Processing
Architecture: *Lisp on the Connection Machine
Principle Investigator: Mr. M A Elmohammed, Syracuse University

This project deals with the implementation of existing models and the modification and extension to existing models, incorporating other aspects the model could serve. I mainly will be dealing with artificial neural network

models which are made of highly parallel rotational elements connected in a pattern that is reminiscent of biological nets. They are of interest primarily because they may be able to emulate the speed and performance of real neural nets using many simple 'slow' computational elements operating in parallel. Thus they offer one possible solution to the problem of obtaining the massive parallelism and computational requirements that are presumed to be required for such problems as signal processing and speech recognition. There has been a great deal of work about the ability of Hopfield's model to serve as a CAM and as a pattern classifier for binary bit patterns (note: the analog side has not been much dealt with). A content-addressable memory retrieves one of (M) stored patterns given an input pattern which is a 'noisy version' of the stored pattern. Basically, I'm focusing on the classification problem for 2 reasons: (1) A CAM is essentially a classifier which outputs the exemplars for selected class instead of an index to that class. In other words, it determines which of (M) exemplar patterns is 'most similar' to a noisy input pattern. (2) Classification is a very fundamental operation that is essential to the important problems of speech, image recognition and signal processing, whether achieved by biological or artificial means. The second part of this study will focus on non-linear signal processing using artificial neural nets with added parameters of prediction and system modelling.

Title:	The Constraint Programming Languages.
Research area:	Logic Programming
Architecture:	Encore Multimax
Principle Investigator:	Mr Peter Crockett, Syracuse University

We believe that the constraint logic programming paradigm provides the correct semantical framework for integrating perceptual and deductive reasoning. Towards this end, we are developing a family of constraint logic programming languages CLP(Mn(F)). Here, (F) is a commutative ring with identity, and $M_n(F)$ is the set of all n-dimensional arrays over (F). A

constraint logic program in CLP(Mn(F) consists of two parts: the first part specifies array operations and the second part is a finite set of definite-clause-like sentences, the antecedent of which has a body and a constraint, while the consequent is an atomic formula. We plan to implement an interpreter for this family of languages. Such an interpreter consists of two parts - one is an ordinary SLD-Resolution theorem prover, while the other is an Array Constraint Solver. The array constraint solver is envisaged as a collection of parallel algorithms that implement various common array operations. The programmer may then specify more complex array operations in terms of these primitive operations and these complex operations could therefore be executed in parallel. Our first goal is to implement the array constraint solver.

Title: 'Real Things in Parallel', a parallel 3-D Natural
Phenomena simulator.
Research area: Computer Graphics
Architecture: Connection Machine and Frame Buffer
Principle Investigator: Gill Fuchs, Syracuse University

Fractal techniques lend themselves beautifully to graphically oriented computer simulations of 'real phenomena'. A recurrent bottle neck that usually arises is due to the serial nature of the implementation—in the computation phase as well as in the display phase. Ergo, a parallel implementation is the one of choice. Furthermore, in image synthesis problems, such as the one at hand, there are several facts that warrant a massively parallel SIMD environment. There is a massive number of display elements (pixels) which correspond to independent calculation nodes (a special case of divide and conquer). Those calculation nodes are relatively independent of each other (in some special cases, completely so) and undertake the same calculation but with slightly differing data. Various phenomena are to be attempted: terrain simulation, cloud formations,

flowers, trees, crystalline growth as well as non-random fractal objects. The first pilot study will consist of a simpler 2D problem to test and define the computational model. The frame buffer of the connection machine will display a Julia set whose iterative complex feedback constant will be read graphically from a Mandelbrot set displayed on a Symbolics monitor. Once a relative performance index is established, a more definite idea of 'real time' simulation capabilities will have been formed.

Title: Analysis of small radomes of arbitrary shape.
Research area: Electromagnetics
Architecture: Fortran on the Alliant
Principle Investigator: Ercument Arvas, Syracuse University

The purpose of this research project is to study the electromagnetic behavior of a small radome of arbitrary shape. The radome is composed of one or more layers of dielectric shells. A small radome has its maximum linear dimension smaller than the wavelength of the operation. The radome may have bends and edges such that the radius of curvature may be much smaller than the wavelength. The radome contains an antenna of arbitrary shape. The antenna is modelled as a metallic body. The radome may also be supported by a perfectly conducting body of arbitrary shape. This body is called 'the reflector'. Given the excitation, we would like to find the field transmitted through the radome which may be backed by the reflector. The problem is solved using the surface equivalence principle to obtain a set of coupled integral equations for the unknown surface equivalent electric and/or magnetic currents. These currents are the sources of the secondary field which exist because of the presence of the radome and/or the reflector. the integral equations are numerically solved using the method of moments.

Title: Benchmarking on the Alliant.
Research area: Computer Assessment
Architecture: Alliant FX/80
Principle Investigator: Mr. Richard K Sato, NCAR

We plan to make a series of runs of the following two programs to test the Alliant: 1) 'Shallow water equations model - this is a fairly simple code which solves a system of partial differential equations on a 2-dimensional rectangular grid. The equations are a primitive but useful approximation to the dynamics of the atmosphere. Second order finite difference approximations are used for the spatial deviatives and leap frog time differencing is used for the time derivative. 2) Climate model - this is a more complex code (20,000 lines) which is a portable version of a 3-dimensional, spectral, global climate model which is used here at the National Center for Atmospheric Research (NCAR). (It is an earlier version of a model which is used by many scientists at NCAR and universities). The model is an out of memory model which therefore makes use of disk files for storage of intermediate data during the mel run. It also starts from an initial data set which is over 4 MBytes.

Title: Efficient Parallel Geometric Operations on Large
Databases
Principle Investigator: Mr. Wm. R Franklin, Rensselaer Polytechnic
Institute

Efficient massively parallel geometric and computer aided design algorithms are the focus of this proposed research. They are often the major compute-bound step in several diverse applications, such as visualization of complex scenes in computer aided design, interference detection in automated assembly, overlay of different maps in cartography, and design rule checking in integrated electronics. Since our techniques are also comparatively easy to

implement, they are the best for applications. Our current results on a sequential machine include processing a complete chip design with 1,819,064 (almost two million) edges. We found all 6,941,110 pairs of intersecting edges in 178 seconds of CPU time. The techniques are extendible to higher level operations. We are now implementing a polyhedron intersection program and a polyhedron visualization program in parallel.

Title: Characterization of Printed Circuits
Principle Investigator: Tapan K Sarkar, Syracuse University

The objective of this research is to develop algorithms to fully characterize multiconductor transmission line networks. The number of networks (possibly non linear) and transmission lines connecting them are arbitrary. The lines are lossy and embedded in multilayered, homogeneous, lossy dielectric media but have uniform cross sections along their lengths. In formulation, we utilize the QUASI-TEM assumption, therefore the lines behavior can be described in terms of their frequency dependent circuit parameters $[R(f)]$, $[L(f)]$, $[C(f)]$, and $[G(f)]$.

Title: A Simultaneous Friction Model with Predicted-Generated Regressors: The Case of Dividend Adjustments
Principle Investigator: Duke Kao, Syracuse University

This study attempts to examine the credibility of dividend signalling in an environment where dynamic dividend adjustments take place constantly. A rational expectations model is proposed to analyze corporate dividend decisions. This model explicitly formulates a formation process of earnings expectations in an efficient market setting. The specification of earnings expectations in the context of efficient markets allows one to isolate the

unexpected future earnings from the expected future earnings changes. Changes in dividend levels are then decomposed into two parts: regular adjustments based on expected earnings and special components dedicated to signal unexpected earnings. This separation of routine partial adjustment from dividend changes for signalling purposes permits a more rigorous test on the effectiveness of using dividends to signal future earnings since dividend changes related to the routine adjustment is under control. Previous studies have failed to draw a distinction between these different dividend changes. Consequently, changes in dividends were all interpreted as a device for signalling future unexpected earnings even though they are possibly a mixture of ordinary adjustments and deliberate changes for signalling. The present study avoids this problem in examining the information content of dividends.

Title: Performance Evaluation of a KEW/Target
Engagement Opportunity Model on the Alliant
Principle Investigator: Ms. Marion Michaud, MITRE Corporation

A comparison of the processing performance achieved for a typical SDI application on various parallel architectures can provide valuable insight into matching these types of algorithms to the most appropriate architectures. A recently developed KEW/Target engagement opportunities model, which is one part of the more global SDI weapon/target allocation problem, is being used as the baseline for such a performance evaluation. Performance measures for this model have already been obtained on RADC's Encore Multimax 320. The aim of this proposal is to obtain similar performance data on the Alliant FX/80 for the same model. Finally, a comparison of the performance and speed-ups achieved on the two machines will be made.

Title: Arcadia Tools to Support Real-time Software Testing
Principle Investigator: Carol H LeDoux, The Aerospace Corporation

The aim of this research project is to design and build prototype tools to support analysis and debugging of real-time and concurrent software written in Ada for embedded systems. These tools include an Ada program execution monitor that collects event streams for post-mortem analysis and an interpreter for matching interval-logic assertions against event streams. Previously we instrumented Ada tasking program to monitor their execution on a single processor and applied interval logic in detecting occurrences of cyclic deadlock and race conditions. We are investigating alternative instrumentation strategies and extending interval-logic analysis to reasoning about the behavior of parallel Ada tasking programs. Because the Encore Multimax offers a parallel Ada environment and access to runtime system instrumentation, we are proposing to monitor task actions in the Multimax runtime system. Key technical challenges are reducing the monitor's interference with program behavior, reasoning about real-time constraints, and allowing for distributed clocks. This project is a component of the Arcadia software development environment program.

Title: Parallel Algorithms for Image Processing and Pattern Recognition
Principle Investigator: Sanjay Ranka, Syracuse University

Many applications in real-time image processing and pattern recognition require computational capability far beyond what can be obtained from the fastest single processor computer available. Hence the usefulness of these applications is limited. Advances in VLSI (Very Large Scale Integration) have provided a cost-effective means of obtaining increased computational power by using parallel computers. The primary thrust of the proposed research is to continue our investigations into the design and

implementation of algorithms for computationally intensive image processing and pattern recognition tasks on parallel computers. We will carry this out in a number of different methods with the overall goal of improving the speed and performance of image processing and pattern recognition systems. We propose to develop efficient computational techniques for some important primitive tasks on different parallel models of computation. We plan to implement our algorithms on the Connection Machine, with 32768 processors, and in the process develop paradigms for effective utilization of commercial SIMD (Single Instruction Stream Multiple Data Stream) architectures like the Connection Machine.

Title: Multi-stage decision processes.

Principle Investigator: John Brule, Syracuse University

There are a large number of interesting problems in such fields as control systems, robotics, signal processing and others that can be classified as multi-stage decision processes. The mathematical and computational demands associated with these processes are significant, and one approach used to solve them in timely fashion is known as dynamic programming. It is proposed to evaluate the effectiveness of the Connection Machine in solving a particular optimization problem using an algorithm based on a dynamic programming solution. The problem to be studied involves a 6-dimensional dynamic system with a figure of merit that is the velocity along a specified trajectory. The state space representation of this problem is quantized into a sufficiently small grid. A parallel machine such as the Connection Machine can be programmed by assigning one or more processors for each data point. The NEWS grid may be utilized and various programming techniques evaluated for optimal solution time. This computational process is widely used in the simulation of physical systems and the results will serve as a benchmark for other solution strategies and an evaluation of the machines.

Title: An extended theory of structural character recognition

Principle Investigator: Dr. James Geller, New Jersey Institute of Technology

The subject matter of the proposed research is real time character recognition. Work in the field of character recognition by computers has followed along three general directions: (1) use of low level features, (2) template matching, and (3) density-quadtrees representation. It is the third technique that we propose to use in this project. The basic idea of density-quadtrees representation is that characters are learned by storing them in a quadtree. A density-quadtrees contains additional information not ordinarily found in a quadtree. The most important item of additional information is their relative pixel density of an area compared to its parent area. In this project we intend to develop an improved density-quadtrees algorithm that adds a normalization step to previous implementations. The nature of the density-quadtrees algorithm allows it to be very gainfully adapted to a Connection Machine implementation; it will make full use of the CM architecture.

Title: Mapping from Abstract Specifications to Parallel Architectures

Principle Investigator: Mr. Richard King, Kestrel Institute

We consider the translation of high-level specification which do not use explicitly parallel control structures into implementations on parallel architectures in categories like a shared memory machine or a transputer or a Connection Machine. Our goal is an automated system that performs such a translation using transformational techniques. We first transform the offered specifications into an intermediate language whose basic elements are the responsibility, grouping responsibilities into processes, the metapredicate, and environment explication. A responsibility can loosely be thought of as the

activity surrounding the computation of the value of a map at one point of its domain; the grouping associates responsibilities into processes, and the metapredicates indicate when resources are to be allocated to responsibilities. Explication of environments allows a logical formula to include information about the manner of a computation as well as about the values to be computed. We believe this language to be well suited for its three roles in our system: as a target language for transformations from first order logic, as a good language for source-to-source transformations, and as a source language for transformation into the target architecture's source language.

Title: Massively Parallel Simulation of Queueing
 Networks on SIMD Machines

Principle Investigator: Mr. Bruce Cota, Syracuse University

Parallel discrete event simulation has been the subject of a great deal of recent work but nearly all of this work has dealt with the MIMD model of computation. Queueing network models are a common subject of discrete event simulation and we believe that the structure of these models will make the use of SIMD computation feasible. We plan to experiment with methods of simulating queueing networks on the connection machine. We intend to use combinations of different methods of parallel simulation in an attempt to exploit the full power of the connection machine.

Title: Parallel Computational Geometry
Principle Investigator: Mr. Frank Dehne, Carleton University

Computational Geometry (CG) methods are powerful algorithm design tools for the efficient solution of a variety of problems occurring in CAD, graphics, image processing, and robotics. To support the design of efficient high performance parallel systems for such applications, I am studying parallel

algorithms for CG and closely related problems on several models of computation; mainly hypercube architectures and processor arrays. During the recent years I have presented a variety of parallel solutions for geometric problems including d-dimensional ECDF searching, 3-dimensional convex hulls, visibility from a point, parallel visibility, separability problems, Voronoi diagrams, computation of the configuration space of a robot, etc. I also studied closely related searching and dictionary problems for processor networks. All results obtained have been published (or have been accepted for publication) in refereed journals or refereed conference proceedings; see enclosed publication list. Most of my methods have so far been implemented and tested by graduate students using the Intel iPSC hypercube. The aim of this research proposal is, to now move these algorithms from this coarse grained medium scale system to a fine grained large scale system like the Connection Machine available at NPAC. The goals are to (1) test the algorithms in a fine grained environment (2) verify, and if necessary modify, current algorithm-analysis methods for fine grained parallel algorithms, and (3) make the methods available also to a wider public in the applied area.

Title: Explicit Methods of Structural Mechanics on Parallel Machines

Principle Investigator: Edward J. Plaskacz, Argonne National Lab

To study the feasibility of finite element analysis on the Connection Machine as well as learning the idiosyncrasies of programming on the Connection machine, a short and simple code in FORTRAN was used as a tool. In deference to the vast volume of FORTRAN code which has been developed over the past 30 years and the infeasibility of 'translating' these codes into a specific language dialect for each innovative computer architecture being developed, the approach I have taken in the execution of my research is to translate the computationally intensive portions of a FORTRAN solid mechanics code into C* and linking it to a FORTRAN driving program which

manages I/O. This approach will undoubtedly capture the interest of organizations using the Connection machine on a time sharing basis in the future for whom extensive code renovation is unfeasible.

Title: Physical Synthesis and Verification of VLSI I
Integrated Circuits
Research area: Design of Very Large Scale Integrated Circuits
Architecture: C/PARIS on CM2 Connection Machine
Principle Investigator: Rob A Rutenbar, Carnegie Mellon University

As the size of Very Large Scale Integrated Circuits (VLSI ICs) increases, the computational demands on the algorithms that synthesize and verify the geometric layout of the required IC masks increases. Despite relentless optimization of sequential algorithms with good serial complexity, the new generation of application-specific ICs with 100,000 to 1,000,000 devices requires enormous amounts of CPU time to design. This project focuses on two parts of this problem: layout synthesis problems based on divide-and-conquer schemes such as graph bipartitioning, and layout verification problems based on computational geometry methods. Both problems are well-suited to massively parallel machines: they require very large data sets, need modest computations on each object in the data set, and exhibit some form of locality of computation. Our goal is to complete parallel algorithms for some of the core layout synthesis and verification steps used in VLSI circuit design, and demonstrate the speedup and quality improvements obtainable with a datawise parallel approach.

Title: Conditional and Parallel Rewriting Laboratory
Principle Investigator: Chilukari Mohan, Syracuse University

Term rewriting systems implement equational theories, with simple operational mechanisms and clear localized declarative semantics. Several implementations have been built on sequential machines, demonstrating the utility of the rewriting paradigm in applications like automated deduction, data type specifications, functional programming, and program verification. We propose to explore the parallelization of various inherently non-deterministic rewriting techniques and applications, which should result in significant speedups. We plan to build fast parallel implementations for rewrite-based applications on existing machines with parallel processing abilities. Another major aspect of the proposed research focuses on conditional rewriting systems, a recent and powerful generalization of term rewriting systems. These consist of sets of oriented equations, each accompanied by an antecedent that needs to be verified before rewriting can occur. Our emphasis is on the implementation of recently developed conditional rewriting mechanisms, with systems containing literals with negation, logical variables and rules ordered by priorities. Implementation on sequential machines will be followed by extending parallel implementations of unconditional rewriting to conditional systems. These conditional and parallel rewriting implementations should be useful for traditional rewriting applications as well as newer problem domains.

Title: Conditional and Parallel Rewriting Laboratory
Principle Investigator: Lan C Li-or, Syracuse University

Term rewriting systems implement equational theories, with simple operational mechanisms and clear localized declarative semantics. Several implementations have been built on sequential machines, demonstrating the utility of the rewriting paradigm in applications like automated deduction,

data type specifications, functional programming, and program verification. We propose to explore the parallelization of various inherently non-deterministic rewriting techniques and applications, which should result in significant speedups. We plan to build fast parallel implementations for rewrite-based applications on existing machines with parallel processing abilities. Another major aspect of the proposed research focuses on conditional rewriting systems, a recent and powerful generalization of term rewriting systems. These consist of sets of oriented equations, each accompanied by an antecedent that needs to be verified before rewriting can occur. Our emphasis is on the implementation of recently developed conditional rewriting mechanisms, with systems containing literals with negation, logical variables and rules ordered by priorities. Implementation on sequential machines will be followed by extending parallel implementations of unconditional rewriting to conditional systems. These conditional and parallel rewriting implementations should be useful for traditional rewriting applications as well as newer problem domains.

Title: Corporate Account: Advanced Simulations of Lasers
Principle Investigator: William J Gehm, Laser Technology Associates, Inc.

Laser Technology Associate Inc. (LTA) is a start-up high-technology company whose major goal is to develop and manufacture advanced solid-state laser devices for the material processing and scientific laser markets. LTA will concentrate its efforts on a novel and relatively new type of laser technology known as a 'slab' laser. Slab lasers have already demonstrated dramatic advantages over the industry standard 'rod' lasers and promise to revolutionize and expand many application areas. In our current development phase, we are creating a suite of sophisticated computer analysis codes to simulate the various physical processes of such laser devices. This is necessary due to the complexity of these lasers and the need to design a device that is optimized for each particular application.

Title: Develop a parallel particle-in-cell code on the connection machine for studying electron beam injections in space.

Principle Investigator: Dr. Chin S Lin, Software Research Institute

The Particle-in-cell (PIC) model, which has been used extensively on supercomputers to simulate fusion and space plasma phenomena, is difficult to develop for the Massively Parallel Processor (MPP) because the model requires indirect indexing in computing electric fields. To overcome the difficulties, a parallel PIC algorithm is developed for MPP by mapping particles in a cell randomly to a row of processors. Because of this mapping, the algorithm needs only the nearest-neighbor communication to sort particles and to collect charge density for each cell. From the cell-charge density, this algorithm then calculates electric fields at the cell by Fast Fourier Transform. The developed PIC code has a speed comparable to that of the vectorized PIC code on CRAY X-MP. The results from simulating the plasma instabilities of a cold electron beam in a hot electron background are presented.

Title: Use of Parallel Computing in predicting the macroscopic properties of two-phase systems.

Principle Investigator: Dr. Ashok S Sangani, Syracuse University

The proposed research is to compute various macroscopic properties of two-phase systems in which the dispersed phase consists of spherical particles. Because of certain efficient numerical schemes that are recently developed by PI and co-workers and because of the availability of efficient computing machines with the facility for parallel computations, it has become possible for the first time to evaluate rigorously various properties of dense systems,

i.e., systems in which the volume fraction of the dispersed phase is large. The results obtained from this research are expected to be a timely contribution to the literature on this general problem of long standing history in a number of branches of theoretical physics and engineering, and may serve as a useful data bank for the future development of more accurate continuum theories and closure schemes in describing the behavior of dispersed systems.

Title: Parallel Implementation of Semantic Networks
Principle Investigator: David F Hille, Syracuse University

This project is an attempt to implement a semantic network on the Connection Machine. A semantic network is a mechanism that stores factual information and attempts to model the associative connections exhibited by humans which make certain items of information accessible from others. Various semantic networks have been implemented on sequential machines. The aim of this project is to explore the possible advantages of parallel processing in the implementation of semantic networks. Specifically, the node space of the semantic network is to be implemented using parallel variables as part of its data structures and the communications and massively parallel processing of the Connection Machine will be used to implement the parser and interpreter of the network. The performance characteristics of a parallel implementation of a semantic network will be compared to the performance characteristics of an equivalent serial implementation.

Title: Digital Enhancement of Remotely Sensed Data for
Mineral Exploration
Principle Investigator: Ashok K Joshi, Syracuse University

Remotely sensed data from satellite platforms have been used extensively for geological exploration since the beginning of the last decade. The blurring of

the image due to modulation of the sensor's optical system, detector and atmosphere often causes degradation of the image. Due to this blurring effect the sharp contacts and edges are not readily visible in the imageries. Photo-interpretation of these satellite images is more commonly used technique and not much emphasis has been paid for digital image processing of remotely sensed data. Available software and hardware have been used to improve the quality of image for manual interpretation but efforts have not been made for quantitative analysis. Also most of the processing for improving the spatial quality of the image has been done in spatial domain using convolution kernels that is not efficient enough to enhance the edges selectively. This type of processing often increases the noise in the image at the expense of resolution improvement. Convolution operators degrade the image quality considerably and that is why, processed images cannot be used for spectral analysis. The restoration for this blurring can be performed in the frequency domain more accurately without increasing the noise in the system restoring the spectral characteristics of the image. However, the conversion of the image from spatial to frequency domain using a Fast Fourier Transform (FFT) involves millions of calculations for a small image of 512x512 pixels. This transformation takes at least 40 minutes or so on available image processors in the campus and that is why it is not very practical to do such processing. The use of fast computers particularly the parallel computations would be able to do the same processing in real time. Processing on Parallel Computers involves Fourier filtering and image restoration techniques to improve the quality of the imageries that can be used for targeting mineral deposits in unexplored regions. The hydrothermal deposits related to the linear features such as faults and shear zones can be located on digitally enhanced remotely sensed data using such techniques that will improve the exploration strategy in mineral projects.

Title: Working With Alopex
Optimization Algorithm
Principle Investigator: Wei Liu, Syracuse University

A research project bases on theoretical investigations and applications of Alopex algorithm which designed to optimize a cost function of many parameters. The optimization procedure is stochastic and iterative. In every iteration all variables that determine the cost function are changed simultaneously by small increments. This feature makes Alopex ideal to implement on parallel computer architectures. The Alopex algorithm has been applied to many practical optimization problems such as traveling salesman problem and crystal growth etc. We now apply the Alopex to the pattern recognition problem. Considering that the number of templates is given, and one input, which is the rotation of one of the templates, is also given. We like to know which template this input represent and what is the rotation angle. Our objective is to find the best match between the input and one of the templates by optimizing the cost function. the process is to make use of the Alopex and rotate all the templates, then the cost function is updated. Since the process is parallel and large arrays are needed to store real objects, using parallel computer can speed up the optimizing process dramatically. The pattern recognition problem we deal with are most three dimension objects, a great number of computations are needed. Computing time is a important factor.

Title: Cellular Automata Applications: Physical Modeling
and Logic-Circuit Simulation
Principle Investigator: Sam Kim, Rensselaer Polytechnic Institute

The goal of our research project is to investigate the application of cellular automata algorithms in two different areas: chemical modeling and logic-circuit operations. The chemical modeling is concentrating on a cellular

automaton model of sublimation above surface. This is accomplished by using a modified probabilistic version of diffusion -limited aggregation. The primary feature of this model under investigation is the temperature-dependence of equilibrium vapor-pressure above the surface. The second application is simulation and compression/expansion of logic-circuits. Compression of the circuit images is desired so that a smaller array can be used for iteration of the simulation algorithm, which will allow the simulation to run faster and use less memory. The actual simulation consists of the standard logic-gates, input and output, and also allows for simulation of faults at arbitrary locations. This project is currently implemented on a 16-processor Sequent.

8. Research Support

Research Support provides a wide range of user services including: Outreach, national access, Research Partners, consulting services, user groups, training, machines for coursework, and lending and reference libraries. The technical research support services are described in this section; the administrative support services are detailed in part two of this report.

8.1 Outreach

All NPAC activities designed to educate the research community about parallel computing, promote its use, and provide active user support have been grouped into a program of Outreach. While there will continue to be special programs developed to specifically accomplish outreach, it is important to note that even routine activities such as consulting and training carry with them components of outreach. Outreach is the consciousness that the actions of NPAC staff have both an immediate and a longer term impact on the public's perception of NPAC and affect the willingness of researchers to use the center. This consciousness is then translated into deliberate action.

The NPAC Outreach program has been active on a small scale, operating primarily at conferences and meetings. Outreach visits have been constrained to date, but the success of our few efforts are most encouraging. We look forward to making and taking more opportunities do more off-site presentations to research groups at research labs and universities.

8.2 National Access and Use of NPAC

Researchers at remote locations, whether across campus or across the country, enjoy the capability of logging on to the NPAC computers and working interactively at reasonable, if not remarkable, speeds. (See *Networking* in section 6) NPAC realizes that the effective support of these researchers presents a special challenge. Their requirements, while basically identical to those of local users, also include the maintenance of the sense and reality of connectivity and community. Every support service is designed to accommodate the remote user as well as the local one. As NPAC has matured and gained a national profile, the proportion of the remote user community has likewise grown. The following institutions are represented among the remote researchers using the NPAC computers.

Cornell University

Carnegie Mellon University

University of Rochester

IBM Corporation, T. J. Watson Research Center

New Mexico State University

Clarkson University

SUNY at Buffalo

SUNY Health Sciences Center

University of Florida

University of Pennsylvania

National Center for Atmospheric Research

California Institute of Technology

Rensselaer Polytechnic Institute

Xerox Corporation, Tarrytown Research Center

Columbia University
Wayne State University
Harvard University
Thinking Machines Corporation

8.3 Research Partners

In recognition of the special requirements of remote users, NPAC designed a distributed support program for providing more immediacy to this community. Called Research Partners, it was predicated on the value of having support available in the form of a person co-located on the same site. Sites wishing to establish an organizational relationship with NPAC to support or promote parallel computing using the NPAC facilities are invited to participate. Membership will be at the institutional or organizational level, rather than at the individual level. The purpose is to build a functional relationship that fosters computing and research, intellectual and actual interaction between on-site users, the NPAC staff and other users.

Research Partners will receive support from NPAC to encourage and facilitate the program's use by their scientists. This may mean on-site documentation for NPAC architectures, preference in NPAC's remote training, opportunities for graduate or faculty exchange programs or internships. Local user groups for Research Partners may get direct, continuing support from NPAC. This is not an inclusive list of the characteristics of the program.

An Advisory Council of representatives from the Partners will be established to help NPAC identify ways in which support for remote users can be improved, and to provide a forum for discussions of frontiers and directions that NPAC might pursue.

8.4 Research Consulting Services

NPAC research consultants provide assistance to researchers using the parallel computers and to those who anticipate such use. Consulting areas include: program strategy design, choosing architectures or algorithms, migrating to a new environment and use of the systems. These consultants do not do custom programming. They are an important interface between the vendors and users, especially when NPAC is engaged in the beta test of new features.

Services are delivered electronically, telephonically and personally. An electronic mailbox is supported, and read, by the consultants on a rotating basis. Tracking of queries and responses is built into a background system, and all transactions are archived.

Telephone consulting is available from 10 am to 5 pm EST, but the phones are generally covered from 8 am to 7 pm EST. Phone requests for help are not currently logged or tracked, but anecdotal evidence suggests that each call consists of multiple 'consult contacts.'

Those researchers who require special or a great deal of assistance will typically arrange a personal consultation.

Research consultants are encouraged to explore the newest features available to NPAC, even anticipating their release to beta test. These staff also teach in the NPAC training workshops and write internal and user documentation.

8.5 User Groups

NPAC has been active in the Alliant User Society, ALLUS, as a participating and a contributing member at the meetings. NPAC has also played a vital role in the maintenance of the Society's on-line software distribution service, known as netlib. NPAC provides the custodial service for this program, and submits quarterly activity reports to Alliant about its use.

Similar to the netlib program is an emerging collaborative effort among Argonne, Naval Research Laboratory and NPAC to provide a similar on-line software exchange service for the Connection Machine. We expect that portions of this library will reside on the NPAC hardware. Once details are final, this service will be announced to users nationwide.

The formation of a Syracuse University "NPAC Parallel Computer User Group" has been planned for several months. A shortage of manpower and scheduling opportunities has prevented its establishment to date. We have maintained some involvement with the SU Supercomputer User Group, which counts NPAC users among its members.

The Association for Computing Machinery supports numerous Special Interest Groups (SIGs), and one of them, University and College Computing Services, SIGUCCS, has created a Supercomputer Task Force. It functions like a user group and places priority on the nationwide dissemination of information about supercomputer and parallel computer centers and available services and support programs. To this end, SIGUCCS is building an on-line information database, accessible to all. NPAC has joined this effort, with two staff attending meetings and contributing information. We expect to continue this relationship.

8.6 Training

NPAC training efforts have expanded during the past year. To date, NPAC has sponsored five separate training seminars and workshops. With each workshop NPAC has been able to assume more responsibility for the logistics, materials presented, and the actual presentations. We have completed a full training workshop for the Alliant FX/80. Workshops for the Encore and CM environments are under development.

The first NPAC-sponsored training was held in October 1987. The topic of the week-long workshop was: An Introduction to Encore Umax Operating System. NPAC's role in this training was limited to arranging logistics and promoting the training to the Syracuse University community. The second training effort was an introduction to the Alliant FX/80 given in April 1988. NPAC handled all of the arrangements for this one-day seminar. For each of these presentations, NPAC relied on the hardware vendor for instructors, materials, as well as the content of the instruction.

The third training session sponsored by NPAC was Umax Parallel Programming offered by Encore Corporation. NPAC took a much more active role in the planning and implementation of this workshop. The vendor supplied NPAC with an instructor and most of the training materials. In addition to handling all of the logistics of the workshop, we also modified the vendor supplied curriculum to more closely match the needs of our research community.

The fourth and fifth workshops were Alliant Training Workshops. These workshops were designed, developed, and delivered by NPAC staff members. Members of Support Services developed an outline for the workshop, members of the technical staff developed original training materials for the workshop, and the entire staff was involved giving presentations and offering support through the workshop laboratory sessions. The curriculum for the second workshop was a refinement of the lectures and laboratory sessions presented during the first sessions.

8.7 Use of NPAC Facilities in Coursework

NPAC offers the use of its facilities for coursework to faculty nationwide. Consistent with NPAC's mission to further the application of parallel computing, we believe in encouraging the next generation of research scientists to become familiar with these architectures as soon as possible. This opportunity has been available for the past two academic semesters. In the

fall of 1988, five courses, for a total of 132 student accounts, at Syracuse University used the center. The courses were:

- VLSI Design Methods
- Advanced VLSI Design Seminar
- Artificial Intelligence for Foreign Policy Analysis
- Research Seminar on Joyce Programming
- Parallel Architectures

In this current spring semester, four courses are offered by Syracuse University, Rensselaer Polytechnic Institute and Carnegie Mellon University are using the NPAC parallel computers in classwork. A total of 82 student accounts were issued for the following courses:

- VLSI Design Methods (SU)
- Engineering Design of Operating Systems (SU)
- Parallel and Distributed Techniques in Artificial Intelligence (RPI)
- Programming Parallel Computers (CMU)

8.8 Publications and Documentation

NPAC has published four versions of its user guide, and for the next release will publish a tandem guide to Unix on NPAC machines. A technical report series has been initiated. A promotional brochure describing the facility and another detailing workshops and training have been produced. NPAC's newsletter, *Parallel Computing News*, is published monthly, and covers topics relevant to NPAC and parallel computing in general.

9. Use of NPAC Facility

9.1 Machine Use

NPAC's Friendly User Program ended in June 1988 when NPAC initiated full operations to establish a base of users whose research focuses on parallel computing and applications related to the use of parallel architectures. The "friendly users" were invited to submit proposals to continue participation as members of the NPAC user community. While this hiatus resulted in a

temporary decrease in machine use, NPAC was now able to focus attention on the target community it was established to support. The Friendly User Program had 760 users when it was phased out in June 1988.

NPAC computers are used to approximately 30% of capacity on the average, with sections of time available for larger projects. The CM2 and Alliant are occasionally subject to full use.

No commercial quality "security systems" are available on the Alliant, Encore, CM2 or CM1. NPAC developed its own system which can detect intrusion. This was done by modifying the source code on each of the operating systems, and is integrated with a single control station. Separate reports are available for intrusion attempts. A formal warning notice has been placed on all machines when users login, enabling NPAC to prosecute those who attempt to break in. Additional software and network security measures are under consideration.

In order to track usage by computer system, the NPAC Systems staff developed techniques to quantify and record demographics such as the number of users and the category of machine use. In the case of the Connection Machine, this involved the actual development of an accounting system as such tools were not available from the vendor. Each month, use is reported back to the Principal Investigators who are responsible for the projects consuming the NPAC resources.

10. Resource Allocation Plan and the Resource Allocation Board

A Resource Allocation Plan was established in conjunction with DARPA and RADC to distribute the unique and limited resources available at NPAC. (A copy of the plan immediately follows this section.) The plan allows for the allocation of 60% of resources to DARPA-approved research projects with the remaining 40% to be distributed by NPAC. Interested researchers and graduate students are asked to submit a resource request form along with a

brief proposal to NPAC when requesting an account. If the researcher is interested in DARPA sponsorship, they are instructed to send the same materials to a stated electronic or US mail address for consideration. NPAC is notified of approved projects directly by DARPA. All other account requests are granted starter accounts immediately, provided they meet minimum requirements, while their research proposal is forwarded to the Resource Allocation Board for peer review. At this time the Resource Allocation Board is comprised of Syracuse University Researchers and Faculty who are involved in computing and/or scientific research. Current plans call for the addition of remote members to the review board beginning in June 1989. Corporate researchers who request accounts are granted accounts outside of the standard process. The terms of their involvement are negotiated with the Assistant Director, Technology Transfer. (See Corporate Partners Program, Section two, for details)

All researchers interested in NPAC facilities are given a Resource Allocation Packet. This packet contains:

- Resource Allocation Policy
- Administration of Research Projects
- Resource Request Procedure
- Format of Request for Resources
- Guide to Estimating Resource Requirements
- Request for Account
- Guidelines for Researchers Submitting Research Proposals
- Guidelines for Graduate Students Submitting Research Proposals
- Guidelines for Faculty Who Wish to Apply for Class Accounts
- Description of Resources: Hardware
- Description of Resources: Software
- Research Support Programs

Administrative Report

Section II

1. Requirements for Success

NPAC must significantly enhance the ease of parallel computing for scientists and engineers engaged in research which demands high performance computing. Success can be measured by: 1) the number of researchers who have used NPAC productively, 2) the quality of their research results, 3) the reduction in the level of effort required by new researchers to begin using NPAC services, 4) the level of user satisfaction.

To achieve success, NPAC must have highly competent and experienced staff providing focused advice and training to researchers; the latest in high performance architectures, operating systems, compilers and applications software; and a demonstrated ability to identify and build solutions for new opportunities in parallel computing.

2. Background on NPAC

NPAC was established in 1987 at Syracuse University to offer an environment dedicated to the study and use of parallel computing.

By providing the latest in hardware and software, including a full suite of common languages, NPAC operates as a test-bed for pioneering research and technological developments in the field of parallel computing. The services and facilities of NPAC are available to qualified researchers in academic, corporate, and government institutions.

When NPAC was founded, it was the first national parallel computing center. NPAC operates under contract to the Rome Air Development Center, Griffiss Air Force Base, Rome, New York, with funding from the Defense Advanced Research Projects Agency, Washington DC.

NPAC supports research in the operation, application and improvement of parallel computers. Using the latest commercially available hardware, with both fine and coarse-grain architectures, NPAC assists scientists and engineers in many areas of work.

NPAC participates in joint research ventures with computer manufacturers to test and evaluate products for further enhancement. Beta-tests of hardware, compilers, operating systems, communications and peripherals are conducted on site. Special allocations of NPAC computing resources are made to researchers for experimentation, stand alone and diagnostic time.

The parallel computers include both massively parallel, single-instruction-multiple-data (SIMD) systems and coarser grained, shared memory, multiple-instruction-multiple-data (MIMD) machines. SIMD architectures are represented by two Connection Machines, supported by high-performance hosts, mass storage facilities and graphical display equipment. MIMD systems include bus-based, distributed cache systems (Encore Multimax), shared cache systems (Alliant FX/80), and special-purpose graphics systems (Stellar GS1000). The computers are robustly configured and are on accelerated upgrade schedules.

In addition to machines, NPAC provides software, documentation and expert consultation to assist researchers using advanced parallel computing techniques. These services and facilities are available nationwide, on the basis of merit review, to qualified researchers in academic, corporate, and government institutions.

While many fields of research are supported at NPAC, areas of particular interest to the center include: artificial intelligence; logic programming; signal processing; image processing; information retrieval; algorithm evaluation; VLSI design; parallel architecture evaluation and measurement.

Nuclear magnetic resonance, molecular dynamics, computational fluid dynamics, real-time imaging, vision and image synthesis are a few of the fields in which scientists are currently using NPAC resources for performing applied research.

Training and education activities include the Invited Lecture Series, workshops, and faculty use of the NPAC computers in teaching. NPAC has

technical and research consulting services and maintains a library of manuals and reference materials. User groups for the various architectures and research interests are also supported.

The NPAC Corporate Program offers companies membership benefits that include: low-cost, low-risk use of NPAC parallel architectures; the Technology Transfer Program; free consulting and training; and participation in other NPAC supported activities.

The University Outreach program encourages use of NPAC parallel computers among the academic research community. Through it, NPAC is building an infrastructure for supporting scientists working remotely from other campuses.

2.1 Architectures at NPAC

At the time of NPAC's move into the Science and Technology Center, the machines included:

Two Connection Machines, a CM-Model 1 and a CM-Model 2, by Thinking Machines Corporation. These SIMD architectures each have 32,768 processors. The CM2 has a potential of more than 5 gigaflops, and a Data Vault storage device with 5 gigabytes of disk space. Two color frame-buffers offer high resolution display from the CM. A Digital Equipment Corporation VAX 8800 and two Symbolics Corporation LISP Workstations serve as CM frontends.

The Multimax 320, by Encore Computer Corporation, is a shared memory bus-based machine with 18 processors, 128 megabytes of shared memory, 36 Mips, 6.4 gigabytes of disk. A less-configured Multimax 320 is also installed; it runs with the Mach operating system—an extension of Berkeley Unix 4.3BSD—which takes advantage of multiprocessor architectures.

The FX/80 from Alliant Computer Systems Corporation, has eight Advanced Computational Elements, six Interactive Processors, 4.5 gigabytes of disk, and

128 megabytes of memory. It is a shared memory parallel computer with fully integrated vector and floating point processing.

The Stellar Graphics Supercomputer Model GS1000 is a four-stream parallel processor that couples computing and graphics in one system. Scientific visualization is possible through the system's 25 Mips/40 Mflops performance combined with 3-D color graphics.

3. Management Team and Staff

NPAC's management team consists of the Director; Associate Director, Core Research; Deputy Director; and Assistant Director for Research Programs, Technology Transfer, Administration, and the Business Manager. This group meets weekly to advise the Director about matters related to the direction of NPAC and to coordinate the efforts of the various program areas. This group makes critical decisions relating to the assessment and planning for new funding and technical opportunities. The current NPAC staff chart is found on the following page.

4. Relationship with Syracuse University

One of NPAC's challenges is trying to meet the demands of a fast-paced technological community while residing in the more deliberately-paced university atmosphere. In many ways, this gives NPAC an ideal base of knowledge and the company of fellow pioneers in computer technology. But we also must try to maintain the delicate balance of respect for the older well established system of operating while trying to accomplish our goals.

Occasionally we met resistance in the face of conflicting styles and priorities. The Director, William Schrader has made a conscious effort this year to reach out to the local SU community to make our colleagues aware of our mission and the challenges we face. On July 25, 1988, he submitted an annual report to Syracuse University to demonstrate the progress NPAC had made since its creation. This report was helpful in bridging the gap of understanding.

Schrader also addressed a well-attended local staff development colloquium to describe NPAC and its mission. Additionally, Schrader and his staff have taken every opportunity to try to foster understanding throughout the local

university community through increased involvement on campus and the use of meetings of selected representatives of key departments. The addition of Elizabeth Schermerhorn as Deputy Director and James Bruel as Assistant Director for Research Programs to the management team has resulted in increased exposure to the local research community and the benefit has been rewarding.

Perhaps the greatest administrative challenge lies in the search for adequate space in which to house our staff and facility. As we have grown rather rapidly to 19 full-time staff members who require space for workstations and peripheral devices as well as a training facility, consultation area and conference room, there has been an ongoing search for additional space to facilitate growth and development. We have grown from two small offices in Machinery Hall to 2370 square feet of office space and about 1000 square feet of machine room space. The offices are fully outfitted with networks, workstations, telephones, and office furniture. The locations, in Machinery Hall and Link Hall, were requested from and graciously provided by Academic Computing Services (ACS), Computer and Network Services (CNS) and the Engineering department.

The newest temporary home for NPAC will be in the Center for Science and Technology, currently nearing completion. Approximately 8000 sq. feet of office space will be loaned by the CASE Center and Chemistry. CNS will lend approximately 3500 sq. feet of machine room space to house all NPAC equipment. Throughout the year, our staff designed this fourth major move in cooperation with Syracuse University's Department of Design and Construction and other future building inhabitants. This space will snugly accommodate 33 people and we will be at full capacity in April. It is a priority for NPAC to acquire space designed to meet our needs and which will support our mission within two years.

4.1 NPAC and programs at Syracuse University

NPAC has significantly contributed to the quality of Syracuse University by participating in an assessment of the University's research environment and

identifying significant strengths in engineering, science, mathematics, business, and the humanities. This process also identified areas for which more fundamental support could be provided, allowing NPAC to deliver its services as productively as possible.

NPAC hired a full-time Assistant Director for Technology Transfer in an effort to promote the University's strengths in computing, especially as they relate to NPAC's goals. While the purpose of the program was to direct corporate attention to NPAC and its strong university setting, this activity prompted general cooperation among faculty, graduate students and support staff. During the year, this position was funded primarily from other sources and therefore could be represented as additional cost sharing by the University.

The great interest in NPAC, generated by these efforts, effected the hiring of a part-time Assistant Director for Research Programs who will promote collaborations and initiate major efforts on campus in fields of interest to NPAC.

Finally, the University has invested in an office of Research Support and a strong, centralized Office of Sponsored Programs to manage both pre- and post-award research efforts. NPAC equipped these offices with micro-computers and communications so the offices would be connected to the LAN and researchers. Electronic mail is now standard between these offices and researchers. Additional NPAC-sponsored projects at Syracuse University are discussed in the Networking section in part one of this report.

5. Administrative research support

(Technical research support is fully described in part one of this report.)

5.1 Invited Lecture Series and Seminar Series

NPAC created an Invited Lecture Series to bring world-renowned leaders in parallel computing research to Syracuse. Lecturers included:

- Joseph Goguen, SRI International, "The Rewrite Rule Machine"

- Kai Hwang, University of Southern California, "Recent Advances in Parallel Processing and Supercomputing"
- Karl Heinz Winkler, Los Alamos National Laboratory, "Data Visualization in Scientific Research"
- James Browne, University of Texas at Austin, "A Practical Unification Model of Parallel Computation and Application to Programming"
- Larry Smarr, University of Illinois at Urbana-Champaign, "Computational Science and Parallel Architectures"
- Geoffrey Fox, California Institute of Technology, "Physical Structure of Concurrent Computers & Parallel Algorithms"
- Daniel Hillis, Thinking Machines Corporation "Using the Connection Machine to Study Evolution or What Happens When a Computer Architect Finally Gets to Use His Own Machine"

A Seminar Series was created to present user results from successful computational research projects on NPAC facilities and elsewhere. Two seminars were held: March 25, 1988, Frank McCabe, Imperial College, London England, "Parallel Architecture for Symbolic Computing", and July 6, 1988, Steve Morton, Oxford Computers, "Intelligent Memory Chips Multiply a Matrix Times a Vector with Unprecedented Power, Flexibility, and Economy"

5.2 Lending Library

NPAC maintains a library of technical documentation, produced and distributed by the computer vendors, available for short-term loan to researchers. The collection contains multiple copies of the most popular and widely used volumes, and pre-release notes for new features, available, for example, under beta test through NPAC.

5.3 Reference Library

To supplement vendor supplied and NPAC generated documentation, a reference library of current subjects and titles has been assembled. General topic areas include architectures, languages, artificial intelligence and parallel and distributed computing. The library is used primarily by NPAC staff, Syracuse University researchers and visiting scientists.

6. Technology Transfer and the Corporate Partners Program

The Corporate Partners Program (CPP) was developed to provide enticing, no-risk, formal association with NPAC for any size business with any specialized interest. The CPP was designed to accommodate all options that corporate interaction with NPAC might present. Those include: regular communication, temporary personnel exchange, on-site training programs or workshops, cash or equipment grants, the most advantageous use of all NPAC resources, access to graduate students with parallel computer experience and training, the exchange of research information, and/or sponsored research.

The CPP requires an annual membership fee of \$10,000. Discounted machine access and free consulting are the two most significant advantages to corporations of the CPP. The *Guide to Corporations Interested in Parallel Computing* lists all of the advantages and free services of the CPP in greater detail. The annual membership provides NPAC with a modest amount of income to underwrite the implementation of the program. Prior to the establishment of the CPP, fees had to be paid for every service requested or required. This was cumbersome to implement and often had the effect of being an impediment to an emerging relationship with a business.

Full implementation of the CPP must be coordinated with all aspects of NPAC technical and administrative operations, the CASE Center, and university-wide operations. The intent is to provide a service that compliments, not competes, with other university centers and policies. In the six months since the structure for the program first took shape, three corporations have become members: General Electric, UTC/Carrier, and Xerox. We have been responding to interest from other businesses, but time is needed to convert their curiosity about parallel computing into commitment to NPAC resources. We believe the CPP provides the right kinds of incentives to significantly add to the tally of corporate partners in the

next 12 months. More detailed information is available in the CPP corporate guide appended to this report.

Discussions are underway with a number of corporations to generate memberships and collaborative research contracts. Advanced discussions include:

- NYNEX has funded an initial contract on multimedia research for approximately \$500,000. A second contract for about \$2,000,000 is expected in 1989. The work will be performed by NPAC, CASE, IST and ECE.
- General Electric, Scientific Applications International, Knowledge Systems Concepts, and NPAC have agreed upon and submitted a joint proposal to RADC and DARPA valued at \$2,550,000.

Other corporations with whom membership discussions are taking place include: Lockheed, Singer-Link, Kodak, Xerox, Motorola, United Technologies, IBM, Digital, NYSERNet, and IIT Research Institute, among others. Special "supplier" relationships were built with Thinking Machines Corporation, Encore, Alliant, Symbolics, and Sun Microsystems. NPAC's Research Partners program covers relationships with other academic institutions and national laboratories.

7. Administrative Summary

Administrative Support is designed to help NPAC meet the financial and contractual requirements of its contract with RADC. It also enables NPAC's technical support staff to reach and assist users.

The development of the communications and office automation system was one of the first projects undertaken by Administrative Support. Systems were chosen to be compatible with our client community and to be continually upgraded. Accordingly, the standard office workstations are Macintosh SEs, Mac II's and SUNs. In addition to serving as terminal access to NPAC computers, these workstations are used for wordprocessing, business

information tracking and reporting, and electronic mail communication. The Macintosh workstations are fully networked internally through the use of Appletalk networks and externally through the campus Ethernet to all available wide area networks (WANs). The SUN workstations are linked through the campus Ethernet, and likewise access the WANs. Workstations are supported with university- and business-standard software such as Microsoft Word for wordprocessing and Microsoft Excel for financial tracking. This common access to software facilitates working in a distributed group. To further enhance compatibility with other workstations, a recently installed system converts DOS 5.25" diskettes to Mac 3.5" diskettes. The monthly newsletter, *Parallel Computing News*, is produced camera-ready with the use of Aldus's Pagemaker, also an industry standard. The Sun workstations support Interleaf for publication of technical articles and the NPAC User Guide. Additional applications are available on both workstation systems to produce graphics, charts and a database. Double Helix II is installed on the Mac and is networked to track proposals and user accounts, maintain our mailing list, and to track personnel and other information. We have formed a representative committee of technical and administrative staff to coordinate the needs of both groups. This committee investigates new applications, updates existing software, maintains awareness of changing staff and users needs, and ensures the continued reliability of our workstation system.

The Business Office has developed a project-tracking interface to the existing University general ledger system. The result of this major project is the ability to provide comprehensive spreadsheets for each contract, including detailed salary, travel, procurement and other information. These spreadsheets are updated monthly and enable NPAC management to see how the financial status of each funding source, fits into the overall picture.

DISTRIBUTION LIST

addresses	number of copies
RADC/COES ATTN: Raymond A. Liuzzi Griffiss AFB NY 13441-5700	2
Northeast Parallel Arch Center Syracuse University Syracuse, New York 13214	5
RADC/DOVL Technical Library Griffiss AFB NY 13441-5700	1
Administrator Defense Technical Info Center DTIC-FDAC Cameron Station Building 5 Alexandria VA 22304-6145	2
Defense Advanced Research Projects Agency 1400 Wilson Blvd Arlington VA 22209-2308	2
RADC/COAC Griffiss AFB NY 13441-5700	1
HQ USAF/SCTT Washington DC 20330-5190	1
SAF/AQSC Pentagon Rm 4D 269 Wash DC 20330	1

Naval Warfare Assessment Center 1
GIDEP Operations Center/Code 3DG
ATTN: E Richards
Corona CA 91720

HQ AFSC/XTH 1
Andrews AFB MD 20334-5000

HQ SAC/SCPT 2
OFFUTT AFB NE 68046

DTESA/RQE 1
ATTN: Mr. Larry G. McManus
Kirtland AFB NM 87117-5000

HQ TAC/DRIY 1
ATTN: Maj. Divine
Langley AFB VA 23665-5575

HQ TAC/DOA 1
Langley AFB VA 23665-5554

ASD/ENEMS 1
Wright-Patterson AFB OH 45433-6503

SM-ALC/MACEA 1
ATTN: Danny McClure
Bldg 237, MASOF
McClellan AFB CA 95652

WRDC/AAAI-4 1
Wright-Patterson AFB OH 45433-6543

WRDC/AAAI-2 1
ATTN: Mr Franklin Hutson
WPAFB OH 45433-6543

AFIT/LDEE 1
Building 642, Area B
Wright-Patterson AFB OH 45433-6583

WRDC/MTEL 1
Wright-Patterson AFB OH 45433

AAMRL/HE 1
Wright-Patterson AFB OH 45433-6573

Air Force Human Resources Lab 1
Technical Documents Center
AFHRL/LRS-TDC
Wright-Patterson AFB OH 45433

AUL/LSE 1
Bldg 1405
Maxwell AFB AL 36112-5564

HQ AFSPACECOM/XRA 1
STINFO Officer
ATTN: Dr. W. R. Matoush
Peterson AFB CO 80914-5001

HQ ATC/TTOI 1
ATTN: Lt Col Killian
Randolph AFB TX 78150-5001

AFLMC/LGY 1
ATTN: Maj. Shaffer
Building 205
Gunter AFS AL 36114-6693

US Army Strategic Def CSSD-IM-PA PO Box 1500 Huntsville AL 35807-3801	1
Ofc of the Chief of Naval Operation ATTN: William J. Cook Navy Electromagnetic Spectrum Mgt Room 5A678, Pentagon (OP-941) Wash DC 20350	1
Commanding Officer Naval Avionics Center Library D/765 Indianapolis IN 46219-2189	1
Commanding Officer Naval Ocean Systems Center Technical Library Code 96428 San Diego CA 92152-5000	1
Cmdr Naval Weapons Center Technical Library/C3431 China Lake CA 93555-6001	1
Superintendent Code 1424 Naval Postgraduate School Monterey CA 93943-5000	1
Space & Naval Warfare Systems Comm Washington DC 20363-5100	1
CDR, U.S. Army Missile Command Redstone Scientific Info Center AMSMI-RD-CS-R/ILL Documents Redstone Arsenal AL 35898-5241	2
Advisory Group on Electron Devices 201 Varick Street, Rm 1140 New York NY 10014	2

Los Alamos National Laboratory
Report Library
MS 5000
Los Alamos NM 87544

1

AEDC Library
Tech Files/MS-100
Arnold AFB TN 37389

1

Commander, USAG
ASQH-PCA-CRL/Tech Lib
Bldg 61801
Ft Huachuca AZ 85613-6000

1

1839 EIG/EIT
Keesler AFB MS 39534-6348

1

AFEWC/ESRI
San Antonio TX 78243-5000

3

485 EIG/EIR
ATTN: S Buzinski
Griffiss AFB NY 13441-6348

1

ESD/XRR
Hanscom AFB MA 01731-5000

1

ESD/AVSE
Hanscom AFB MA 01731-5000

1

ESD/SZM
Hanscom AFB MA 01731-5000

1

SEI JPO 1
ATTN: Major Charles J. Ryan
Carnegie Mellon University
Pittsburgh PA 15213-3890

Director NSA/CSS 1
T513/TDL
ATTN: D W Marjarum
Fort Meade MD 20755-6000

Director NSA/CSS 1
W157
9800 Savage Road
Fort Meade MD 21055-6000

NSA 1
ATTN: D. Alley
Div X911
9800 Savage Road
Ft Meade MD 20755-6000

Director 1
NSA/CSS
W11 DEFSMAC
ATTN: Mr. Mark E. Clesh
Fort George G. Meade MD 20755-6000

Director 1
NSA/CSS R12
ATTN: Mr. Dennis Heinbuch
9800 Savage Road
Fort George G. Meade MD 20755-6000

DoD 1
R31
9800 Savage Road
Ft. Meade MD 20755-6000

DIRNSA 1
R509
9800 Savage Road
Ft Meade MD 20775

Director 1
NSA/CSS
R08
Fort George G. Meade MD 20755-6000

DOD Computer Center
C/TIC
9800 Savage Road
Fort George G. Meade MD 20755-6000

1

DARPA/ISTO
1400 Wilson Blvd
Arlington VA 22209-2308

1



MISSION of Rome Air Development Center

RADC plans and executes research, development, test and selected acquisition programs in support of Command, Control, Communications and Intelligence (C³I) activities. Technical and engineering support within areas of competence is provided to ESD Program Offices (POs) and other ESD elements to perform effective acquisition of C³I systems. The areas of technical competence include communications, command and control, battle management information processing, surveillance sensors, intelligence data collection and handling, solid state sciences, electromagnetics, and propagation, and electronic reliability/maintainability and compatibility.